

Posterior Differential Regularization with f -divergence for Improving Model Robustness

Hao Cheng¹, Xiaodong Liu¹, Lis Pereira², Yaoliang Yu³, Jianfeng Gao¹

¹Microsoft Research ²Ochanomizu University ³University of Waterloo & Vector Institute

{chehao, xiaodl, jfgao}@microsoft.com

{kanashiro.pereira}@ocha.ac.jp

{yaoliang.yu}@uwaterloo.ca

Abstract

We address the problem of enhancing model robustness through regularization. Specifically, we focus on methods that regularize the model posterior difference between clean and noisy inputs. Theoretically, we provide a connection of two recent methods, Jacobian Regularization and Virtual Adversarial Training, under this framework. Additionally, we generalize the posterior differential regularization to the family of f -divergences and characterize the overall framework in terms of Jacobian matrix. Empirically, we compare those regularizations and standard BERT training on a diverse set of tasks to provide a comprehensive profile of their effect on model generalization. For both fully supervised and semi-supervised settings, we show that regularizing the posterior difference with f -divergence can result in well-improved model robustness. In particular, with a proper f -divergence, a BERT-base model can achieve comparable generalization as its BERT-large counterpart for in-domain, adversarial and domain shift scenarios, indicating the great potential of the proposed framework for enhancing NLP model robustness.¹

1 Introduction

Although recent neural network based models have achieved great success in a wide range of natural language processing (NLP) tasks, these models may still suffer catastrophic degradation in out-of-domain generalization to datasets with domain shift or adversarial scenarios (Nie et al., 2019; Hsieh et al., 2019). For example, large-scale pretrained neural language models (Devlin et al., 2019; Liu et al., 2019) have better generalization, but experience performance reduction caused by domain shift (Hendrycks et al., 2020). Textual entailment models trained on MNLI (Williams et al., 2018) have picked up superficial cues focusing on either the

¹Code is available at <https://github.com/hao-cheng/f-divergence>.




Figure 1: BERT-based question answering (QA) and natural language inference (NLI) model robustness towards domain shift and adversarial attack. The numbers in percentage indicate the corresponding relative performance drop between in-domain and out-of-domain test sets.

presence of certain keywords (Gururangan et al., 2018) or whether similar words are mentioned in the sentence pairs (McCoy et al., 2019). Jia and Liang (2017) have also shown that SQuAD models are very easily distracted by irrelevant sentences that contain many question words regardless of context, despite of their human-performance on in-domain data. As shown in Figure 1, three BERT-based (Devlin et al., 2019) models perform well for in-domain evaluation data, but transfer poorly to out-of-domain datasets with domain shift or adversarial attack, *i.e.* more than 25% relative performance reduction.

Achieving good generalizations towards datasets with domain shift has been a long-standing goal of domain adaptation. Various methods (Blitzer et al., 2007; Daumé III, 2007) have been developed for training models to learn effectively from both in-domain (source) and out-of-domain (target) datasets. Additionally, the recent discovery on the prevalence of data biases (Gururangan et al., 2018; Jia and Liang, 2017; McCoy et al., 2019), unintended correlations between input and output learned by statistical models, ignites the develop-

ment of model debiasing techniques (Clark et al., 2019; He et al., 2019). These methods leverage discovered data biases to improve model generalization over adversarial datasets (Jia and Liang, 2017; McCoy et al., 2019) designed to fool naively trained models. Instead of relying on knowledge of the target dataset, we focus on task-agnostic training techniques for enhancing model robustness with access to only in-domain data.

Motivated by recent success of adversarial training in computer vision (Madry et al., 2018; Goodfellow et al., 2014) and NLP (Zhu et al., 2020; Jiang et al., 2020; Cheng et al., 2019; Wang et al., 2019), we investigate a regularization framework, which directly regularizes the model posterior difference for clean and noisy inputs, as a means to enhance the model robustness. Here, we first provide a theoretical connection of two recent methods under this framework, *i.e.* Virtual Adversarial Training (VAT) (Miyato et al., 2018) and Jacobian Regularization (JR) (Sokolić et al., 2017). In addition, we propose to generalize VAT and random perturbation training (RPT) (Miyato et al., 2018) with a family of probability distribution metrics, f -divergences, and characterize their connection with JR.

Given that large-scale pretrained neural language models have demonstrated their superior generalization for downstream NLP tasks under both matched (Devlin et al., 2019; Liu et al., 2019) and mismatched evaluations (Hendrycks et al., 2020), we systematically study the regularization framework using BERT (Devlin et al., 2019) on a diverse set of tasks in terms of both in-domain and out-of-domain generalization. Specifically, we use representative datasets (Socher et al., 2013; Williams et al., 2018; Rajpurkar et al., 2016, 2018) from sentiment analysis, textual entailment and question answering (QA) for in-domain training and evaluation. In order to assess the resulting model generalization over domain shift and adversarial attack, we then consider out-of-domain datasets (Tsatsaronis et al., 2015; Maas et al., 2011) and challenge adversarial datasets (Jia and Liang, 2017; McCoy et al., 2019) in a zero-shot learning fashion.

Our experiments show that regularizing the posterior difference for clean and noise inputs is very effective in improving model generalization under both supervised and semi-supervised learning settings. Based on our theoretical analysis, both VAT and RPT variants, unlike JR, incorporate model confidence for adaptive regularization, which leads

to consistently better empirical robustness over BERT with the standard fine-tuning. Furthermore, we find that different f -divergences lead to different generalization behaviors for in-domain, domain shift and adversarial settings. In our study, VAT with symmetric divergence achieve better generalization for in-domain and domain shift cases, while VAT with asymmetric divergence achieve more robustness toward adversarial attack. More importantly, we show that a BERT-base model trained with a proper f -divergence can perform comparably to its corresponding BERT-large counterpart. It is also worth noting that VAT with symmetric divergence lead to improved data efficiency, *i.e.* achieving comparable in-domain performance as fully-supervised models with only 50% labelled data. This further illustrates the great potential of the proposed general regularization framework for the semi-supervised setting.

Our main contributions are summarized as follows: 1) we generalize the posterior differential regularization framework to the family of f -divergences and provide additional divergence functions with different characteristics for regularization; 2) based on our framework, we analyze the family of regularization methods and show the theoretical connection of recently proposed methods, JR, VAT and RPT, in terms of their regularization effect on the input-output Jacobian matrix; 3) We provide a comprehensive profile of different regularization methods over a diverse set of NLP tasks and experimental insight into which f -divergence is more suitable for improving NLP model robustness under both supervised and semi-supervised settings.

2 Posterior Differential Regularization

In this section, we first introduce the regularization framework that penalizes the difference of posterior between clean and noisy inputs. Based on this framework, we set up the basic notions of two recent methods, *i.e.* VAT and JR, and show their theoretical connection. Finally, we generalize the posterior differential regularization with any function from the family of f -divergences and characterize their local smoothness promoting in terms of Jacobian matrix. In the following, we use $f_{\Theta}(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ to denote the posterior function which is a neural network parameterized by Θ that maps the input $\mathbf{x} \in \mathbb{R}^n$ to the output probability space with m discrete classes.

Both adversarial learning (Goodfellow et al., 2014) and the posterior difference regularization aim at making the model more robust. Adversarial learning focuses on minimizing the following objective

$$\min_{\Theta} \max_{\|\epsilon\| \leq c} L(f_{\Theta}(\mathbf{x} + \epsilon), y), \quad (1)$$

where L is the cross-entropy loss, $\epsilon \in \mathbb{R}^n$ is a random vector bounded in a norm by c , a small positive constant and y is the target label. Instead, the posterior differential regularization directly promotes the model local smoothness, *e.g.* stabilizing the model posterior distribution towards small input perturbations. Typically, it is in the form of

$$\min_{\Theta} L(f_{\Theta}(\mathbf{x}), y) + \alpha R(f_{\Theta}(\mathbf{x}), f_{\Theta}(\hat{\mathbf{x}})), \quad (2)$$

where $\hat{\mathbf{x}} = \mathbf{x} + \epsilon$, R is a regularization term penalizing the model instability, and α is a hyperparameter for balancing the classification loss and the regularization term. As we can see, posterior differential regularization is a task-agnostic method which makes it applicable to semi-supervised, self-supervised and unsupervised learning. For simplicity, we will use f and R to denote f_{Θ} and $R(f_{\Theta}(\mathbf{x}), f_{\Theta}(\hat{\mathbf{x}}))$, respectively.

Jacobian Regularization: A recent regularization approach to stabilize the model is Jacobian regularization (Sokolić et al., 2017; Li et al., 2016). Specifically, using the input-output Jacobian matrix, $J = J_f(\mathbf{x}) \in \mathbb{R}^{m \times n}$, we can get the first-order Taylor approximation

$$f(\hat{\mathbf{x}}) = f(\mathbf{x} + \epsilon) = f(\mathbf{x}) + J_f(\mathbf{x})\epsilon. \quad (3)$$

In order to reduce the overall model sensitivity to the input perturbation, Sokolić et al. (2017) propose to directly regularize the Frobenius norm of the input-output Jacobian matrix so that

$$\begin{aligned} \|f(\hat{\mathbf{x}}) - f(\mathbf{x})\|_2^2 &= \|J\epsilon\|_2^2 = \epsilon^T J^T J \epsilon \\ &\leq \|\epsilon\|_2^2 \|J\|_{sp}^2 \leq \|\epsilon\|_2^2 \|J\|_F^2, \end{aligned}$$

where $\|\cdot\|_2$ is the L_2 norm, $\|\cdot\|_{sp}$ is the spectral norm, and $\|J\|_F^2 = \text{tr}(J^T J)$ is the Frobenius norm of the Jacobian matrix with tr as the trace operator. In other words, by letting $R = \|J\|_F^2$, the L_2 difference between clean and noisy inputs is thus being effectively regularized.

Virtual Adversarial Training: Motivated by the adversarial learning objective used in (Goodfellow

et al., 2014), Miyato et al. (2018) introduce a regularized objective to enhance the model robustness towards small input perturbations

$$\min[L(y, \hat{y}) + \alpha \underbrace{\max_{\|\epsilon\| \leq c} \text{KL}(\hat{y}, f(\mathbf{x} + \epsilon))}_R], \quad (4)$$

where KL is the well-known *Kullback-Leibler divergence*, and $\hat{y} = f(\mathbf{x})$. Based on the above definition, VAT essentially regularizes the KL-based worst-case posterior difference between the clean and noisy input using an inner loop to search for the most adversarial direction.

Although sharing with JR the same spirit of encouraging the local smoothness of the model, JR and VAT are not fully theoretically connected. In what follows, we use a simple approach to draw the theoretical connection between these two methods.

Connection between VAT and JR: Here, we show that VAT and JR can be directly related through the definition of induced matrix norm. Specifically, the matrix norm of the Jacobian matrix is

$$\|J\| = \sup_{\|\nu\|=1} \|J\nu\|, \quad (5)$$

where the matrix norm on the left side is induced by the corresponding vector norm on the right side. It is easy to show that

$$\begin{aligned} c^2 \|J\|_{sp}^2 &\approx \sup_{\|\epsilon\|_2=c} \|J\epsilon\|_2^2 \\ &\leq \sup_{\|\epsilon\|_2=c} \|f(\hat{\mathbf{x}}) - f(\mathbf{x})\|_1^2 \\ &\leq 2 \sup_{\|\epsilon\|_2=c} \text{KL}(f(\mathbf{x}), f(\hat{\mathbf{x}})), \end{aligned}$$

and the last inequality is attained based on Pinsker's inequality. Therefore, the VAT regularization provides an upper bound for the spectral norm of the Jacobian matrix. Although a similar attempt to relate JR and VAT has been first explored in (Abbas et al., 2016), we provide a simple and comprehensive connection. Specifically, both VAT and JR regularize the upper bound of the spectral norm of the Jacobian matrix.

Posterior Differential Regularization with f -divergence: Although both VAT and JR have been successful in improving model robustness, they are both special cases of regularizing the model posterior difference between the clean and noisy inputs. One natural question is whether we can

use other probability distribution metrics for regularization and characterize them in terms of Jacobian matrix. In the following, we extend the posterior difference regularization with the family of f -divergences (Csiszár and Shields, 2004). Furthermore, we show that posterior differential regularization with all f -divergences results in an adaptive variant of JR which incorporates model confidence.

First, let's define the f -divergence for measuring the posterior difference between the clean and noisy inputs, *i.e.*

$$D_g(f(\hat{\mathbf{x}}), f(\mathbf{x})) = \sum_i f_i(\mathbf{x}) g\left(\frac{f_i(\hat{\mathbf{x}})}{f_i(\mathbf{x})}\right), \quad (6)$$

where the *generator function* $g : \mathbb{R}_+ \mapsto \mathbb{R}$ is a convex and lower-semicontinuous function satisfying $g(1) = 0$, $\hat{\mathbf{x}} = \mathbf{x} + \epsilon$ and f_i indicates the i -th element of vector f . Different choices of g lead to several popular divergences, *e.g.* KL, squared Hellinger and Jensen-Shannon divergence. Based on this, it is easy to show that the corresponding second order approximation is

$$D_g(f(\hat{\mathbf{x}}), f(\mathbf{x})) \approx \frac{g''(1)}{2} \epsilon^T J^T \text{diag}\left(\frac{1}{f}\right) J \epsilon, \quad (7)$$

where J is the input-output Jacobian of f , and $\text{diag}\left(\frac{1}{f}\right)$ is a diagonal matrix with elements equal to $\frac{1}{f}$ (See Appendix A for full derivation).

Compared with the Frobenius norm of Jacobian matrix $\|J\|_F^2$, Equation 7 can be seen as a weighted version of JR where each row is rescaled by the model confidence f_i for the corresponding class. In other words, it is close to JR for more confident classes, whereas for uncertain classes it allows less Jacobian variance. Additionally, although $g''(1)$ is a constant once the generator function is selected, various f -divergences can lead to different approximations which might result in task-dependent benefits. Therefore, different from KL-based VAT or its sampling alternative without the inner search for the most adversarial direction as proposed in (Miyato et al., 2018), we generalize the posterior differential regularization with the family of f -divergences and show that they all provide an approximation to a variant of JR which adapts the regularization based on model confidence.

3 Model

Given its superior performance over a wide range of NLP tasks, we focus on exploring different train-

ing techniques using BERT (Devlin et al., 2019). We first describe the BERT representations used for all tasks considered in this paper. Then, two variants of task-specific BERT-based models are introduced: 1) the sentence-level classifier for textual entailment and sentiment analysis, and 2) the extractive QA model. Specifically, we focus on different ways of encoding input text and building task-specific layers using BERT representations.

BERT Representation: For all tasks considered in this work, an input text sequence is divided into subword units $w_t, t = 1, \dots, T$. The tokenized input sequence is then transformed into embeddings, $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathbb{R}^n$, through a token encoder, which combines a token embedding, a (token) position embedding and a segment embedding (*i.e.*, which text span the token belongs to) by element-wise summation. The embedding layer is used as the input to multiple transformer layers (Vaswani et al., 2017) to generate the contextual representations, $\mathbf{h}_1, \dots, \mathbf{h}_T \in \mathbb{R}^d$, which are the hidden states of the last layer of the BERT model. For all regularizations, we sample noise vectors $\epsilon_1, \dots, \epsilon_T$ from $\mathcal{N}(0, I)$, and normalize each vector into L2 unit vector. The noise input is then constructed by adding the normalized noise vector to the token embeddings, *i.e.* $\mathbf{x}_1 + c\epsilon_1, \dots, \mathbf{x}_T + c\epsilon_T$. Here, we fix $c = 1e-3$ in this paper.

Sentence-level Classifier: Following the standard setup of BERT-based textual entailment model (Devlin et al., 2019), a pair of premise and hypothesis is converted into an input sequence in the form of "[CLS] *premise* [SEP] *hypothesis* [SEP]". Here, [CLS] is a special token indicating the start of the whole sequence and [SEP] is another special token for separating the two sentences. For sentiment analysis, a single sentence is converted to the form of "[CLS] *sentence* [SEP]".

For both classification tasks, the task-specific layer only takes the first hidden vector $\mathbf{h}_{[CLS]}$ produced by BERT, corresponding to the [CLS] token. Then, the probability of class k is

$$P(k|w_1, \dots, w_T) \propto W_k^C \mathbf{h}_{[CLS]}, \quad (8)$$

where $W^C \in \mathbb{R}^{m \times d}$ is the learnable parameter, the subscript k indicates the k -th row of the matrix, and the bias term is left out for simplicity. For standard BERT training, the log-likelihood based on Equation 8 is used. For regularized models, the regularization term is added to stabilize the class probability change with regard to the input noise.

Extractive QA Model: For extractive QA, the probability space outcomes consist of token positions of answer spans. Given a pair of question q and a passage p in the form of "[CLS] *question* [SEP] *passage* [SEP]", the BERT encoder produces contextualized representations for all tokens in the input. Specifically, for each token position t in p , the final hidden vector $\mathbf{h}_t \in \mathbb{R}^d$ is used as the contextualized token embedding, where d is the vector dimension.

The span-begin score is computed as $s_b(i) = \mathbf{w}_b^T \mathbf{h}_i$ using a weight vector $\mathbf{w}_b \in \mathbb{R}^d$. The probability for a start position i is

$$P_b(i) = \frac{\exp(s_b(i))}{Z_b}, \quad (9)$$

where Z_b is the normalizing factor computed by normalizing over \mathcal{I} (the set of all possible positions in the passage), *i.e.* $Z_b = \sum_{i \in \mathcal{I}} \exp(s_b(i))$. The span-end score $s_e(j)$, the probability $P_e(j)$ for an end position j , and the normalizing factor Z_e are defined in the same way. The probability of an answer span (i, j) is

$$P(i, j) = P_b(i)P_e(j) = \frac{\exp(s_b(i) + s_e(j))}{Z_b Z_e}.$$

Maximizing the log-likelihood of the above equation is equivalent to maximizing the log probabilities for the correct start and end position, respectively. For regularized models, given it is computationally expensive to enumerate all possible spans, we apply two separate regularization terms for the start and end position probabilities, respectively.

4 Experiments

In this section, we apply the regularization methods discussed so far to BERT and evaluate their performance on the model robustness. Specifically, we consider two types of posterior regularization with f -divergences. In addition to a VAT-like regularization with an inner search for the most adversarial direction, following (Miyato et al., 2018), we also evaluate the random perturbation training (RPT) with the family of f -divergences which only uses randomly sampled noise for regularization. In this work, we focus on three representative f -divergences, *i.e.* KL, *squared Hellinger* (SHL) and *Jensen-Shannon divergence* (JSD).

Dataset: All the datasets used in this paper are summarized in Table 1. We consider three tasks, *i.e.* QA, textual entailment, and sentiment analysis, where the last two are sentence classification

tasks. Following the literature, we report the exact match (EM) and F1 scores for QA datasets and classification accuracy for textual entailment and sentiment analysis. For model training, we use MNLI (Williams et al., 2018) and SST-2 (Socher et al., 2013) and SQuAD v1.1/v2.0 (Rajpurkar et al., 2016, 2018), respectively. The corresponding development set is used for evaluating the in-domain generalization.

To evaluate the out-of-domain generalization with domain shift, we use the BioAQS dataset (Tsatsaronis et al., 2015) from MRQA (Fisch et al., 2019) and the IMDB dataset (Maas et al., 2011). Unlike SQuAD which is based on Wikipedia, BioAQS is a biomedical QA dataset constructed on PubMed articles. Compared with SST-2 containing pithy expert reviews (Socher et al., 2013), IMDB includes lengthy movie reviews from non-experts (Maas et al., 2011). We directly apply the QA model trained on SQuAD v2.0 and the sentiment classifier trained on SST-2 to BioAQS and IMDB, respectively.

To evaluate the model robustness towards adversarial attack, we use two challenging adversarial datasets, *i.e.* Adversarial SQuAD (Jia and Liang, 2017) and HANS (McCoy et al., 2019) for evaluating QA model trained on SQuAD v1.1 and the textual entailment model trained on MNLI, respectively. The Adversarial SQuAD is constructed based on SQuAD v1.1 (Rajpurkar et al., 2016) by adding distracting sentences that have high overlap with the question and contain plausible answer candidates. Naively trained models tend to exploit the word overlap with the given question and thus are fooled by those distracting sentences (Jia and Liang, 2017). The HANS dataset is built using three heuristics to ensure that the hypothesis sentence only contains words from the premise sentence (McCoy et al., 2019). Similarly, standard training results in models failing catastrophically, even for BERT.

Implementation: We follow the default setting used for fine-tuning the uncased BERT base model (Devlin et al., 2019). We select the learning rate from $\{3e-5, 4e-5\}$ for QA models and $\{2e-5, 3e-5\}$ for classification models. For both tasks, we tune the number of training epochs in $\{2, 3, 4, 5\}$. In addition, we search regularization weight in $\{0.001, 0.01, 0.1\}$ for JR, and $\{1, 4, 10\}$ for VAT and RPT. We use the in-domain dev set for validation and select the best model based on F1

Task	Training	Metrics	Evaluation	
			Domain Shift	Adversarial Attack
Question Answering	SQuAD v1.1	F1/Exact Match (EM)	N/A	Adversarial SQuAD
Question Answering	SQuAD v2.0	F1/Exact Match (EM)	BioASQ	N/A
Textual Entailment	MNLI	Accuracy (Acc)	N/A	HANS
Sentiment Analysis	SST-2	Accuracy (Acc)	IMDB	N/A

Table 1: Summary of datasets and their corresponding evaluation purpose and metrics. N/A indicates the trained model is not evaluated for the corresponding generalization. For full data statistics, see [Appendix B](#).

Method	v1.1	v2.0	MNLI	SST-2	Avg Δ
	F1	F1	Acc	Acc	
BERT _{base}	88.5	76.5	84.5	92.3	0.0
JR	88.5	74.2	84.5	92.0	-0.7
RPT _{KL}	89.3	77.6	84.8	93.0	+0.7
RPT _{SHL}	89.7	77.6	85.2	93.0	+0.9
RPT _{JSD}	89.8	77.5	85.4	93.5	+1.1
VAT _{KL}	89.9	78.9	85.3	93.1	+1.4
VAT _{SHL}	90.2	79.2	85.5	93.2	+1.6
VAT _{JSD}	90.2	79.3	85.7	93.6	+1.8
BERT _{large}	90.9	81.8	86.3	93.5	+2.7

Table 2: Comparison of different training techniques for in-domain generalization on SQuAD (v1.1 and v2.0), MNLI and SST-2 dev sets. BERT_{base} and BERT_{large} stand for BERT-base and BERT-large model with standard training respectively. The corresponding best performing BERT-base model is in bold. The Avg Δ column is the corresponding average score difference with BERT_{base}.

for QA tasks and accuracy for classification tasks.

In-domain: In this part, we focus on comparing the in-domain performance of different training methods. In other words, each model is trained on the training set and evaluated on the corresponding matched development set. The experiment is summarized in Table 2. In general, JR performs similarly to the standard BERT training with an exception case for SQuAD v2.0. This is probably because JR uniformly regularizes the Jacobian matrix, which is particularly problematic for QA task with unanswerable questions. Both RPT and VAT with different f -divergences achieve significant improvement over standard training for all four datasets, especially on SQuAD v2.0. The results suggest incorporating the model confidence into regularization can achieve better in-domain

Method	BioASQ	IMDB	Avg Δ
	F1/EM	Acc	
BERT _{base}	57.1/41.7	87.7	0.0
JR	60.8/46.0	87.4	+1.7
RPT _{KL}	59.2/43.6	88.7	+1.6
RPT _{SHL}	60.0/44.8	88.7	+1.9
RPT _{JSD}	58.3/43.2	88.3	+0.9
VAT _{KL}	60.1/45.7	86.7	+1.0
VAT _{SHL}	60.7/45.9	87.4	+1.7
VAT _{JSD}	61.8/47.0	88.3	+2.6
BERT _{large}	63.5/49.5	88.3	+3.5

Table 3: Domain shift evaluation of different training techniques on BioASQ and IMDB. BERT_{base} and BERT_{large} stand for BERT-base and BERT-large model with standard training respectively. The corresponding best performing BERT-base model is in bold. The Avg Δ column is the corresponding average score difference with BERT_{base}. For BioASQ, F1 score is used for computing the average score. EM stands for the exact match score.

generalization. Consistent with findings in (Miyato et al., 2018), by searching for the most adversarial perturbation direction, VAT variants achieve the largest boost for in-domain generalization. Moreover, we find that both RPT and VAT with SHL and JSD provides additional improvement over their corresponding counterpart with KL which suggests the benefit of using alternative f -divergences for posterior difference regularization. Lastly, by selecting the proper divergence, the performance gap between the BERT-base and BERT-large model is dramatically narrowed which indicates the advantage of applying posterior difference regularization with f -divergences on top of powerful text representations.

Domain Shift: In this part, we compare the performance of models trained using different techniques

Method	AddSent F1/EM	AddOneSent F1/EM	HANS Acc	Avg Δ
BERT _{base}	54.0/48.9	64.8/59.0	58.4	0.0
JR	55.5/49.3	65.3/58.4	62.6	+2.1
RPT _{KL}	55.2/49.2	66.4/59.8	59.9	+1.4
RPT _{SHL}	57.8/52.2	68.9/62.7	60.0	+3.1
RPT _{JSD}	56.0/49.9	67.1/61.1	61.5	+2.5
VAT _{KL}	60.6/55.0	69.4/62.9	67.7	+6.8
VAT _{SHL}	60.1/ 55.1	69.1/ 63.1	63.2	+5.1
VAT _{JSD}	58.5/53.2	67.4/61.4	64.6	+4.4
BERT _{large}	59.3/54.3	69.1/63.6	67.9	+6.4

Table 4: Adversarial evaluation of different training techniques on Adversarial SQuAD (AddSent and AddOneSent) and HANS. BERT_{base} and BERT_{large} stand for BERT-base and BERT-large model with standard training respectively. The corresponding best performing BERT-base model is in bold. The Avg Δ column is the corresponding average score difference with BERT_{base}. For Adversarial SQuAD, F1 score is used for computing the average score. EM stands for the exact match score.

on datasets with domain shift, *e.g.* different topic or style. Specifically, we apply the QA models trained on SQuAD v2.0 to the BioAQS version from MRQA (Fisch et al., 2019). Similarly, we apply the sentiment analysis model trained on SST-2 to the IMDB test set. The results are summarized in Table 3. Comparing Table 3 with Table 2, all methods suffer a noticeable performance drop for both QA and sentiment analysis when evaluated on test sets with domain shift. Moreover, we observe more significant performance drop for the QA setting because the biomedical domain differs significantly from the Wiki domain in topic and style, resulting in a larger domain shift between the training and test QA datasets. Consistent with findings in (Hendrycks et al., 2020), the in-domain performance is not predictive of the domain shift generalization. Further, the performance of JR is not stable, with improvement on BioASQ but worse performance on IMDB. Models trained with all three RPT variants result in consistent improvement over standard training on both out-of-domain datasets, suggesting that random perturbation is particularly effective in enhancing model robustness towards domain shift. In particular, all RPT variants achieve comparable out-of-domain gen-

eralization on IMDB as BERT-Large. Although all VAT variants achieve decent improvement on BioASQ, neither VAT_{KL} nor VAT_{SHL} generalize so well to IMDB. This illustrates the importance of selecting a proper divergence for VAT style regularization. In other words, domain-dependent search for the most adversarial direction with either KL or SHL might be suboptimal for model generalization over domain shift.

Adversarial Attack: Here, we evaluate different training techniques on adversarial attack scenarios, where datasets are intentionally constructed to fool naively trained models. Specifically, we evaluate the QA models trained with SQuAD v1.1 and the textual entailment models learned on MNLI using the Adversarial SQuAD and the HANS datasets, respectively. Table 4 summarizes the evaluation results of model robustness towards adversarial attacks with different training methods. For both subsets (AddSent and AddOneSent) from Adversarial SQuAD and HANS, all regularization methods improve over standard BERT training. In this case, models trained with VAT variants demonstrate stronger resilience towards learning superficial cues from data. Specifically, VAT with KL achieves the largest improvement on both settings which indicates that an asymmetrical divergence might be more effective in avoiding learning data biases. Although better text representations derived from BERT-Large are still more robust against adversarial attack than the base version, this gap can be effectively reduced by regularizing the posterior difference with f -divergences. Compared with the recent debiasing method proposed in (Clark et al., 2019) that requires the knowledge of existing data bias, VAT variants can be an effective task-agnostic debiasing approach with better in-domain performance and comparable improvement for adversarial settings.

Semi-supervised Learning: One advantage of regularization methods is their compatibility with semi-supervised learning. Given JR is not very effective for the fully-supervised learning, we focus on evaluating RPT and VAT with f -divergences under the semi-supervised setting. Specifically, we use the two sentence classification datasets, MNLI and SST-2, for training. We hold out 50% of the label information for the training data. For standard BERT training, only the labelled part is used. For both RPT and VAT variants, the rest unlabelled data is also included for training. Both the cross en-

	SST-2	IMDB	MNLI	HANS
BERT _{full}	92.3	87.7	84.5	58.4
BERT _{base}	91.2	86.3	82.7	51.5
RPT _{KL}	91.3	87.0	83.6	53.8
RPT _{SHL}	91.9	86.6	83.8	53.3
RPT _{JSD}	91.7	86.5	83.7	51.8
VAT _{KL}	92.1	86.3	83.1	54.3
VAT _{SHL}	92.4	86.5	84.4	51.8
VAT _{JSD}	92.2	86.6	84.1	52.6

Table 5: Comparison of different methods with semi-supervised learning in classification accuracy. Except BERT_{full}, all others are trained with 50% labelled data.

ropy loss and the regularization term are optimized for the labelled samples, whereas only the regularization term is used for unlabelled ones. Similar to the fully supervised setting, the models trained on MNLI is applied to HANS for evaluating the model robustness towards adversarial attack, and the models using SST-2 are applied to IMDB to assess the model performance under domain shift. Results are summarized in Table 5.

Compared with the fully supervised setting, all methods get lower classification accuracy across the board. Both RPT and VAT variants again improve over standard training for both in-domain and out-of-domain evaluations. It is worth mentioning that both SHL and JSD based VAT models trained with 50% labelled data on SST-2 and MNLI are on par with the corresponding standard BERT training with the full training set which illustrates the advantage of choosing a proper f -divergence for the semi-supervised setting. With only half labelled data, Both RPT and VAT suffer a large drop on HANS and produce almost random predictions, indicating the complimentary benefits of data diversity. We also further reduce the amount of labelled training data and observe the same trend where regularizing with different f -divergences can lead to improved data efficiency. This demonstrates the potential of posterior differential regularization for NLP with low-resource scenarios.

5 Related Work

With the goal of developing more robust NLP models, a line of recent work has been devoted to identifying various kinds of superficial patterns learned by high-performance models over many

popular datasets (Gururangan et al., 2018; Jia and Liang, 2017; McCoy et al., 2019). The prevalence of data *biases* over popular datasets poses a real challenge of accurately estimating the model capacity for practical applications, because a closed dataset evaluation usually inflates the model performance. This concern about dataset biases has led researchers to develop new diagnostic datasets (Jia and Liang, 2017; McCoy et al., 2019; Nie et al., 2019) and training techniques (Clark et al., 2019; He et al., 2019) to overcome those discovered biases. Recent debiasing methods (Clark et al., 2019; He et al., 2019) require learning multiple models to access known data biases for the target dataset. Moreover, they achieve more robust out-of-domain generalization at the price of in-domain performance degradation. In contrast, we focus on the task-agnostic robust learning framework for enhancing model robustness and empirically show that regularization approaches under this framework can result in superior in-domain and out-of-domain generalization.

Training with noise has been a very popular approach for enhancing model robustness. The dropout is a widely-used approach in deep learning to improve model generalization (Srivastava et al., 2014). For adversarial learning methods, the main theme is reducing the model sensitivity toward small input perturbations (Goodfellow et al., 2014; Madry et al., 2018), which has been recently applied to both fine-tuning (Jiang et al., 2020; Pereira et al., 2020; Zhu et al., 2020; Li and Qiu, 2020) and pre-training (Liu et al., 2020). However, models trained with adversarial learning are found to have at-odd generalization (Tsipras et al., 2019; Zhang et al., 2019). Our work studies learning methods with the goal of regularizing the model posterior difference of clean and noisy inputs. We show that compared with the standard BERT training, the proposed posterior differential regularization with f -divergence lead to better NLP model robustness.

6 Conclusion

In this paper, we investigate methods regularizing the posterior difference between the clean and noisy inputs for improving model generalization for both in-domain and out-of-domain settings. Specifically, we present theoretical analyses of three methods under this framework, *i.e.* VAT, JR, and RPT. We further extend both VAT and PRT to the family of f -divergences and theoretically characterize them in terms of Jacobian matrix. We also demonstrate

their effectiveness in enhancing model robustness over a diverse set of NLP tasks under both fully-supervised and semi-supervised scenarios.

For future work, it is interesting to explore posterior differential regularization methods for weakly-supervised learning, such as relation extraction and QA with distant supervision.

Acknowledgement

We would like to thank the anonymous reviewers for valuable suggestions. Yaoliang Yu thanks NSERC for funding support.

References

- Mudassar Abbas, J. Kivinen, and T. Raiko. 2016. Understanding regularization by virtual adversarial training, ladder networks and others. In *ICLR 2016-Workshop Track*.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. [Biographies, Bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 440–447, Prague, Czech Republic. Association for Computational Linguistics.
- Yong Cheng, Lu Jiang, and Wolfgang Macherey. 2019. [Robust neural machine translation with doubly adversarial inputs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4324–4333, Florence, Italy. Association for Computational Linguistics.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. [Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4069–4082, Hong Kong, China. Association for Computational Linguistics.
- I. Csiszár and P. Shields. 2004. [Information theory and statistics: A tutorial](#). *Foundations and Trends in Communications and Information Theory*, 1(4):417–528.
- Hal Daumé III. 2007. [Frustratingly easy domain adaptation](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 256–263, Prague, Czech Republic. 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.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. [MRQA 2019 shared task: Evaluating generalization in reading comprehension](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 1–13, Hong Kong, China. Association for Computational Linguistics.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- He He, Sheng Zha, and Haohan Wang. 2019. Unlearn dataset bias for natural language inference by fitting the residual. In *EMNLP Workshop on DeepLo*.
- Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song. 2020. [Pretrained transformers improve out-of-distribution robustness](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2744–2751, Online. Association for Computational Linguistics.
- Yu-Lun Hsieh, Minhao Cheng, Da-Cheng Juan, Wei Wei, Wen-Lian Hsu, and Cho-Jui Hsieh. 2019. On the robustness of self-attentive models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1520–1529.
- Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2020. [SMART: Robust and efficient fine-tuning for pretrained natural language models through principled regularized optimization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2177–2190, Online. Association for Computational Linguistics.
- Linyang Li and Xipeng Qiu. 2020. Textat: Adversarial training for natural language understanding with token-level perturbation. *arXiv preprint arXiv:2004.14543*.

- Yitong Li, Trevor Cohn, and Timothy Baldwin. 2016. [Learning robust representations of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1979–1985, Austin, Texas. Association for Computational Linguistics.
- Xiaodong Liu, Hao Cheng, Pengcheng He, Weizhu Chen, Yu Wang, Hoifung Poon, and Jianfeng Gao. 2020. Adversarial training for large neural language models. *arXiv preprint arXiv:2004.08994*.
- 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*.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. [Towards deep learning models resistant to adversarial attacks](#). In *International Conference on Learning Representations*.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2019. Adversarial nli: A new benchmark for natural language understanding. *arXiv preprint arXiv:1910.14599*.
- Lis Pereira, Xiaodong Liu, Fei Cheng, Masayuki Asahara, and Ichiro Kobayashi. 2020. Adversarial training for commonsense inference. *arXiv preprint arXiv:2005.08156*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- J. Sokolić, R. Giryes, G. Sapiro, and M. R. D. Rodrigues. 2017. Robust large margin deep neural networks. *IEEE Transactions on Signal Processing*, 65(16):4265–4280.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15(56):1929–1958.
- George Tsatsaronis, Georgios Balikas, Prodromos Malakasiotis, Ioannis Partalas, Matthias Zschunke, Michael R Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, Yannis Almirantis, John Pavlopoulos, Nicolas Baskiotis, Patrick Gallinari, Thierry Artières, Axel-Cyrille Ngonga Ngomo, Norman Heino, Eric Gaussier, Liliana Barrio-Alvers, Michael Schroeder, Ion Androutsopoulos, and Georgios Paliouras. 2015. [An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition](#). *BMC Bioinformatics*, 16(1):138.
- Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. 2019. [Robustness may be at odds with accuracy](#). In *International Conference on Learning Representations*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Dilin Wang, Chengyue Gong, and Qiang Liu. 2019. Improving neural language modeling via adversarial training. In *International Conference on Machine Learning*, pages 6555–6565.
- 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.

Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. 2019. [Theoretically principled trade-off between robustness and accuracy](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 7472–7482. PMLR.

Chen Zhu, Yu Cheng, Zhe Gan, Siqu Sun, Tom Goldstein, and Jingjing Liu. 2020. [FreeLB: Enhanced adversarial training for natural language understanding](#). In *International Conference on Learning Representations*.

A Second-order Approximation of Posterior Differential Regularization with f -divergence

Based on Equation 6, it is easy to show that the first and second derivatives of D with regard to ϵ are

$$D' = \sum_i g'(\mathbf{r}_i) f'_i(\hat{\mathbf{x}}),$$

$$D'' = \sum_i \left[g''(\mathbf{r}_i) \frac{f'_i(\hat{\mathbf{x}}) f_i'^T(\hat{\mathbf{x}})}{f_i(\mathbf{x})} + g'(\mathbf{r}_i) f_i''(\hat{\mathbf{x}}) \right],$$

where $\mathbf{r}_i = \frac{f_i(\hat{\mathbf{x}})}{f_i(\mathbf{x})}$. Let $D(\epsilon) = D(f(\hat{\mathbf{x}}), f(\mathbf{x}))$. When evaluating at 0, the second order approximation of the f -divergence is

$$D(\epsilon) \approx D(0) + D'^T(0)\epsilon + \frac{\epsilon^T D''(0)\epsilon}{2}. \quad (10)$$

Given that

$$D(0) = \sum_i g(1) f_i = 0,$$

$$D'(0) = \sum_i g'(1) f'_i = g'(1) \left(\sum_i f_i \right)' = 0,$$

$$\sum_i f_i'' = \left(\sum_i f_i \right)'' = 0$$

the second order approximation can be simplified as

$$(10) = \frac{g''(1)}{2} \epsilon^T \left(\sum_i \frac{f'_i f_i'^T}{f_i} \right) \epsilon + \frac{g'(1)}{2} \epsilon^T \left(\sum_i f_i'' \right) \epsilon$$

$$= \frac{g''(1)}{2} \sum_i \frac{\epsilon^T f'_i f_i'^T \epsilon}{f_i}$$

$$= \frac{g''(1)}{2} \epsilon^T J^T \text{diag} \left(\frac{1}{f} \right) J \epsilon, \quad (11)$$

where J is the input-output Jacobian of f , and $\text{diag} \left(\frac{1}{f} \right)$ is a diagonal matrix with its diagonal elements equal to the vector $\frac{1}{f}$.

B Dataset Statistics

The statistics of all datasets used in this paper is summarized in Table 6. The top four datasets in Table 6 are used for training and the matched dev sets are used for both validation and in-domain evaluation. The remaining datasets are used only for

Corpus	Train	Dev	Test	Metrics
SST-2	67k	872	-	Acc
MNLI	393k	20k	-	Acc
SQuAD v1	87.6k	10.5k	-	EM/F1
SQuAD v2	130.3k	11.9k	-	EM/F1
IMDB	-	-	25k	Acc
BioASQ	-	1504	-	EM/F1
AddSent	-	-	3560	EM/F1
AddOneSent	-	-	1787	EM/F1
HANS	-	-	30k	Acc

Table 6: Experiment dataset summary. The top four datasets are used for training and the matched dev sets are used for both validation and in-domain evaluation. The rest five datasets are used only for evaluation purpose.

evaluation purpose. Among the evaluation datasets, other than BioASQ (Tsatsaronis et al., 2015) from MRQA where the development set is used for evaluation, the corresponding test set is used for evaluating model out-of-domain generalization.