

Adversarial Regularization as Stackelberg Game: An Unrolled Optimization Approach

Simiao Zuo^{†*}, Chen Liang[†], Haoming Jiang[□], Xiaodong Liu[◇], Pengcheng He[◇],
Jianfeng Gao[◇], Weizhu Chen[◇] and Tuo Zhao[†]

[†]Georgia Institute of Technology [□]Amazon [◇]Microsoft
{simiaozuo, cliang73}@gatech.edu, jhaoming@amazon.com
{xiaodl, Pengcheng.H, jfgao, wzchen}@microsoft.com,
tourzhao@gatech.edu

Abstract

Adversarial regularization has been shown to improve the generalization performance of deep learning models in various natural language processing tasks. Existing works usually formulate the method as a zero-sum game, which is solved by alternating gradient descent/ascent algorithms. Such a formulation treats the adversarial and the defending players equally, which is undesirable because only the defending player contributes to the generalization performance. To address this issue, we propose Stackelberg Adversarial Regularization (SALT), which formulates adversarial regularization as a Stackelberg game. This formulation induces a competition between a leader and a follower, where the follower generates perturbations, and the leader trains the model subject to the perturbations. Different from conventional approaches, in SALT, the leader is in an advantageous position. When the leader moves, it recognizes the strategy of the follower and takes the anticipated follower’s outcomes into consideration. Such a leader’s advantage enables us to improve the model fitting to the unperturbed data. The leader’s strategic information is captured by the Stackelberg gradient, which is obtained using an unrolling algorithm. Our experimental results on a set of machine translation and natural language understanding tasks show that SALT outperforms existing adversarial regularization baselines across all tasks. Our code is publicly available.

1 Introduction

Adversarial regularization (Miyato et al., 2017) has been shown to improve the generalization performance of deep learning models in various natural language processing (NLP) tasks, such as language modeling (Wang et al., 2019b), machine translation (Sato et al., 2019), natural language understanding (Jiang et al., 2020), and reading comprehen-

sion (Jia and Liang, 2017). However, even though significant progress has been made, the power of adversarial regularization is not fully harnessed.

Conventional adversarial regularization is formulated as a zero-sum game (a min-max optimization problem), where two players seek to minimize/maximize their utility functions. In this formulation, an adversarial player composes perturbations, and a defending player solves for the model parameters subject to the perturbed inputs. Existing algorithms find the equilibrium of this zero-sum game using alternating gradient descent/ascent (Madry et al., 2018). For example, in a classification problem, the adversarial player first generates the input perturbations by running projected gradient ascent to maximize a loss function, and then the defending player updates the model using gradient descent, trying to decrease the classification error. Notice that in this case, neither of the players know the strategy of its competitor, i.e., the model does not know how the perturbations are generated, and vice versa. In other words, the two players are of the same priority, and either one of them can be advantageous in the game. It is possible that the adversarial player generates over-strong perturbations that hinder generalization of the model.

To resolve this issue, we grant the defending player (i.e., the model) a higher priority than the adversarial player by letting the defender recognize its competitor’s strategy, such that it is advantageous in the game. Consequently, we propose Stackelberg Adversarial Regularization (SALT), where we formulate adversarial regularization as a Stackelberg game (Von Stackelberg, 2010). The concept arises from economics, where two firms are competing in a market, and one of the them is in the leading position by acknowledging the opponent’s strategy. In Stackelberg adversarial regularization, a leader solves for the model parameters, and a follower generates input perturbations. The leader procures its advantage by considering what the best response

*Corresponding author.

of the follower is, i.e., how will the follower respond after observing the leader’s decision. Then, the leader minimizes its loss, anticipating the predicted response of the follower.

The SALT framework identifies the interaction between the leader and the follower by treating the follower’s strategy (i.e., the input perturbations) as an operator of the leader’s decision (i.e., the model parameters). Then we can solve for the model parameters using gradient descent. One caveat is that computing the gradient term, which we call the Stackelberg gradient, requires differentiating the interaction operator. To rigorously define this operator, recall that the follower can be approximately solved using gradient ascent. We can treat the perturbations in each iteration as an operator of the model parameters, and the interaction operator is then the composition of such update-induced operators. Correspondingly, the Stackelberg gradient is obtained by differentiating through these updates. This procedure is referred to as unrolling (Pearlmutter and Siskind, 2008), and the only computational overhead caused by it is computing Hessian vector products. As a result, when applying the finite difference method, computing the Stackelberg gradient requires two backpropagation and an extra $O(d)$ complexity operation, where d is the embedding dimension. Therefore, the unrolling algorithm computes the Stackelberg gradient without causing much computational overhead.

We conduct experiments on neural machine translation (NMT) and natural language understanding (NLU) tasks. For the NMT tasks, we experiment on four low-resource and one rich-resource datasets. SALT improves upon existing adversarial regularization algorithms by notable margins, especially on low-resource datasets, where it achieves up to 2 BLEU score improvements. To test performance on NLU tasks, we evaluate SALT on the GLUE (Wang et al., 2019a) benchmark. SALT outperforms state-of-the-art models, such as BERT (Devlin et al., 2019), FreeAT (Shafahi et al., 2019), FreeLB (Zhu et al., 2019), and SMART (Jiang et al., 2020). We build SALT on the BERT-base architecture, and we achieve an average score of 84.5 on the GLUE development set, which is at least 0.7 higher than existing methods. Moreover, even though we adapt SALT to BERT-base, the performance is noticeably higher than the vanilla BERT-large model (84.5 vs. 84.0).

The unrolling procedure was first proposed

for auto-differentiation (Pearlmutter and Siskind, 2008), and later applied in various context, such as hyper-parameter optimization (Maclaurin et al., 2015; Finn et al., 2017), meta-learning (Andrychowicz et al., 2016), and Generative Adversarial Networks (Metz et al., 2017). To the best of our knowledge, we are the first to apply the unrolling technique to adversarial regularization to improve generalization performance.

We summarize our contributions as the following: (1) We propose SALT, which employs a Stackelberg game formulation of adversarial regularization. (2) We use an unrolling algorithm to find the equilibrium of the Stackelberg game. (3) Extensive experiments on NMT and NLU tasks verify the efficacy of our method.

Notation. We use $df(x)/dx$ to denote the gradient of f with respect to x . We use $\partial f(x, y)/\partial x$ to denote the partial derivative of f with respect to x . For a d -dimensional vector v , its ℓ_2 norm is defined as $\|v\|_2 = (\sum_{i=1}^d v_i^2)^{1/2}$, and its ℓ_∞ norm is defined as $\|v\|_\infty = \max_{1 \leq i \leq d} |v_i|$.

2 Background and Related Works

◇ **Neural machine translation** has achieved superior empirical performance (Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017). We focus on the Transformer architecture (Vaswani et al., 2017), which integrates the attention mechanism in an encoder-decoder structure. The encoder in a Transformer model first maps a source sentence into an embedding space, then the embeddings are fed into several encoding layers to generate hidden representations, where each of the encoding layers contains a self-attention mechanism and a feed-forward neural network (FFN). After which the Transformer decoder layers, each contains a self-attention, a encoder-decoder attention, and a FFN, decode the hidden representations.

◇ **Adversarial training** was originally proposed for training adversarial robust classifiers in image classification (Szegedy et al., 2014; Goodfellow et al., 2015; Madry et al., 2018). The idea is to synthesize strong adversarial samples, and the classifier is trained to be robust to them. Theoretical understanding (Li et al., 2019) about adversarial training and various algorithms to generate the adversarial samples, such as learning-to-learn (Jiang et al., 2021), are proposed. Besides computer vision, adversarial training can also benefit reinforce-

ment learning (Shen et al., 2020). Different from the above fields, in NLP, the goal of adversarial training is to build models that generalize well on the unperturbed test data. Note that robustness and generalization are different concepts. Recent works (Raghunathan et al., 2020; Min et al., 2020) showed that adversarial training can hurt generalization performance, i.e., accuracy on clean data. As such, adversarial training needs to be treated with great caution. Therefore, in NLP, this technique requires refined tuning of, for example, the training algorithm and the perturbation strength.

◊ **Fine-tuning pre-trained language models** (Peters et al., 2018; Devlin et al., 2019; Radford et al., 2019; Liu et al., 2019b; He et al., 2020) is state-of-the-art for natural language understanding tasks such as the GLUE (Wang et al., 2019a) benchmark. Recently, there are works that use adversarial pre-training (Liu et al., 2020a) and adversarial-regularized fine-tuning methods such as SMART (Jiang et al., 2020), FreeLB (Zhu et al., 2019), and FreeAT (Shafahi et al., 2019) to improve model generalization and robustness (Cheng et al., 2021).

3 Method

Natural language inputs are discrete symbols (e.g., words), instead of continuous ones. Therefore, a common approach to generate perturbations is to learn continuous embeddings of the inputs and operate on the embedding space (Miyato et al., 2017; Clark et al., 2018; Sato et al., 2018, 2019; Stutz et al., 2019). Let $f(x, \theta)$ be our model, where x is the input embedding, and θ is the model parameter. Further let y be the ground-truth output corresponding to x . For example, in NMT, f is a sequence-to-sequence model, x is the embedding of the source sentence, and y is the target sentence. In classification tasks, f is a classifier, x is the input sentence/document embedding, and y is the label. In both of these cases, the model is trained by minimizing the empirical risk over the training data, i.e.,

$$\min_{\theta} \mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(f(x_i, \theta), y_i).$$

Here $\{(x_i, y_i)\}_{i=1}^n$ is our dataset, and ℓ is a task-specific loss function, e.g., cross-entropy loss.

3.1 Adversarial Regularization

Adversarial Regularization (Miyato et al., 2017) is a regularization technique that encourages smooth-

ness of the model outputs around each input data point. Concretely, we define an adversarial regularizer for non-regression tasks as

$$\ell_v(x, \delta, \theta) = \text{KL}(f(x, \theta) \parallel f(x + \delta, \theta)),$$

$$\text{where } \text{KL}(P \parallel Q) = \sum_k p_k \log \frac{p_k}{q_k}.$$

Here $\text{KL}(\cdot \parallel \cdot)$ is the Kullback–Leibler (KL) divergence, δ is the perturbation corresponding to x , and $f(\cdot, \theta)$ is the prediction probability simplex given model parameters θ . In regression tasks, the model output $f(\cdot, \theta)$ is a scalar, and the adversarial regularizer is defined as

$$\ell_v(x, \delta, \theta) = (f(x, \theta) - f(x + \delta, \theta))^2.$$

Then the training objective is

$$\min_{\theta} \mathcal{L}(\theta) + \frac{\alpha}{n} \sum_{i=1}^n \max_{\|\delta_i\| \leq \epsilon} \ell_v(x_i, \delta_i, \theta), \quad (1)$$

where α is a tuning parameter, ϵ is a pre-defined perturbation strength, and $\|\cdot\|$ is either the ℓ_2 norm or the ℓ_∞ norm.

The min and max problems are solved using alternating gradient descent/ascent. We first generate the perturbations δ by solving the maximization problem using several steps of projected gradient ascent, and then we update the model parameters θ with gradient descent, subject to the perturbed inputs. More details are deferred to Appendix A.

One major drawback of the zero-sum game formulation (Eq. 1) is that it fails to consider the interaction between the perturbations δ and the model parameters θ . This is problematic because a small change in δ may lead to a significant change in θ , which renders the optimization ill-conditioned. Thus, the model is susceptible to underfitting and generalize poorly on unperturbed test data.

3.2 Adversarial Regularization as Stackelberg Game

We formulate adversarial regularization as a Stackelberg game (Von Stackelberg, 2010):

$$\begin{aligned} \min_{\theta} \mathcal{F}(\theta) &= \mathcal{L}(\theta) + \frac{\alpha}{n} \sum_{i=1}^n \ell_v(x_i, \delta_i^K(\theta), \theta), \\ \text{s.t. } \delta_i^K(\theta) &= U^K \circ U^{K-1} \circ \dots \circ U^1(\delta_i^0). \end{aligned} \quad (2)$$

Here “ \circ ” denotes operator composition, i.e., $f \circ g(\cdot) = f(g(\cdot))$. Following conventions, in this

Stackelberg game, we call the optimization problem in Eq. 2 the leader. Further, the follower in Eq. 2 is described using a equality constraint. Note that U^K is the follower’s K -step composite strategy, which is the composition of K one-step strategies $\{U^k\}_{k=1}^K$. In practice, K is usually small. This is because in NLP, we target for generalization, instead of robustness, and choosing a small K prevents over-strong adversaries.

In Eq. 2, U^k s are the follower’s one-step strategies, and we call them update operators, e.g., U^1 updates δ^0 to δ^1 using pre-selected algorithms. For example, projected gradient ascent can be applied as the update procedure, that is,

$$\begin{aligned} \delta^k(\theta) &= U^k(\delta^{k-1}(\theta)) \\ &= \Pi_{\|\cdot\| \leq \epsilon} \left(\delta^{k-1}(\theta) + \eta \frac{\partial \ell_v(x, \delta^{k-1}(\theta), \theta)}{\partial \delta^{k-1}(\theta)} \right) \\ &\text{for } k = 1, \dots, K, \end{aligned} \quad (3)$$

where $\delta^0 \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ is a initial random perturbation drawn from a normal distribution with variance $\sigma^2 \mathbf{I}$, η is a pre-defined step size, and Π denotes projection to the ℓ_2 -ball or the ℓ_∞ -ball.

To model how the follower will react to a leader’s decision θ , we consider the function $\delta^K(\theta)$. Then, adversarial training can be viewed solely in terms of the leader decision θ .

We highlight that in our formulation, the leader knows the strategy, instead of only the outcome, of the follower. This information is captured by the Stackelberg gradient $d\mathcal{F}(\theta)/d\theta$, defined as the following:

$$\begin{aligned} \frac{d\mathcal{F}(\theta)}{d\theta} &= \frac{d\ell(f(x, \theta), y)}{d\theta} + \alpha \frac{d\ell_v(x, \delta^K(\theta), \theta)}{d\theta} \\ &= \underbrace{\frac{d\ell(f(x, \theta), y)}{d\theta} + \alpha \frac{\partial \ell_v(x, \delta^K, \theta)}{\partial \theta}}_{\text{leader}} \\ &\quad + \alpha \underbrace{\frac{\partial \ell_v(x, \delta^K(\theta), \theta)}{\partial \delta^K(\theta)} \frac{d\delta^K(\theta)}{d\theta}}_{\text{leader-follower interaction}}. \end{aligned} \quad (4)$$

The underlying idea behind Eq. 4¹ is that given a leader’s decision θ , we take the follower’s strategy into account (i.e., the “leader-follower interaction” term) and find a direction along which the

¹The second term in “leader” is written as $\partial \ell_v(x, \delta^K, \theta)/\partial \theta$, instead of $\partial \ell_v(x, \delta^K(\theta), \theta)/\partial \theta$. This is because the partial derivative of θ is only taken w.r.t. the third argument in $\ell_v(x, \delta^K, \theta)$. We drop the θ in $\delta^K(\theta)$ to avoid causing any confusion.

Algorithm 1: Stackelberg Adversarial Regularization with Unrolled Optimization.

Input: \mathcal{D} : dataset; T : total number of training epochs; σ^2 : variance of initial perturbations; K : number of unrolling steps; Optimizer: optimizer to update θ .

Initialize: model parameters θ ;

for $t = 1, \dots, T$ **do**

for $(x, y) \in \mathcal{D}$ **do**

 Initialize $\delta^0 \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$;

for $k = 1, \dots, K$ **do**

 Compute δ^k using Eq. 3;

 Compute $d\delta^k(\theta)/d\theta$ using Eq. 6;

end

 Compute $d\mathcal{F}(\theta)/d\theta$ based on $d\delta^K(\theta)/d\theta$ using Eq. 4;

$\theta \leftarrow \text{Optimizer}(d\mathcal{F}(\theta)/d\theta)$;

end

end

Output: θ

leader’s loss decreases the most. Then we update θ in that direction. Note that the gradient used in standard adversarial training (Eq. 1) only contains the “leader” term, such that the “leader-follower interaction” is not taken into account.

3.3 SALT: Stackelberg Adversarial Regularization

We propose to use an unrolling method (Pearlmutter and Siskind, 2008) to compute the Stackelberg gradient (Eq. 4). The general idea is that since the interaction operator is defined as the composition of the $\{U^k\}$ operators, all of which are known, we can directly compute the derivative of $\delta^K(\theta)$ with respect to θ . Concretely, we first run a forward iteration to update δ , and then we differentiate through this update to acquire the Stackelberg gradient.

Note that the updates of δ can take any form, such as projected gradient ascent in Eq. 3, or more complicated alternatives like Adam (Kingma and Ba, 2015). For notation simplicity, we denote $\Delta(x, \delta^{k-1}(\theta), \theta) = \delta^k(\theta) - \delta^{k-1}(\theta)$. Accordingly, Eq. 3 can be rewritten as

$$\delta^k(\theta) = \delta^{k-1}(\theta) + \Delta(x, \delta^{k-1}(\theta), \theta). \quad (5)$$

The most expensive part in computing the Stackelberg gradient (Eq. 4) is to calculate $d\delta^K(\theta)/d\theta$,

which involves differentiating through the composition form of the follower’s strategy:

$$\frac{d\delta^k(\theta)}{d\theta} = \frac{d\delta^{k-1}(\theta)}{d\theta} + \frac{\partial\Delta(x, \delta^{k-1}, \theta)}{\partial\theta} + \frac{\partial\Delta(x, \delta^{k-1}(\theta), \theta)}{\partial\delta^{k-1}(\theta)} \frac{d\delta^{k-1}(\theta)}{d\theta}$$

for $k = 1, \dots, K$. (6)

We can compute Eq. 6 efficiently using deep learning libraries, such as PyTorch (Paszke et al., 2019). Notice that $\Delta(x, \delta^{k-1}(\theta), \theta)$ already contains the first order derivative with respect to the perturbations. Therefore, the term $\partial\Delta(x, \delta^{k-1}(\theta), \theta)/\partial\delta^{k-1}(\theta)$ contains the Hessian of $\delta^{k-1}(\theta)$. As a result, in Eq. 4, the most expensive operation is the Hessian vector product (Hvp). Using the finite difference method, computing Hvp only requires two backpropagation and an extra $O(d)$ complexity operation. This indicates that in comparison with conventional adversarial training, SALT does not introduce significant computational overhead. The training algorithm is summarized in Algorithm 1.

4 Experiments

In all the experiments, we use *PyTorch*² (Paszke et al., 2019) as the backend. All the experiments are conducted on NVIDIA V100 32GB GPUs. We use the *Higher* package³ (Grefenstette et al., 2019) to implement the proposed algorithm.

4.1 Baselines

We adopt several baselines in the experiments.

- ◊ *Transformer* (Vaswani et al., 2017) achieves superior performance in neural machine translation.
- ◊ *BERT* (Devlin et al., 2019) is a pre-trained language model that exhibits outstanding performance after fine-tuned on downstream NLU tasks.
- ◊ *Adversarial training* (Adv, Sato et al. 2019) in NMT can improve models’ generalization by training the model to defend against adversarial attacks.
- ◊ *FreeAT* (Shafahi et al., 2019) enables “free” adversarial training by recycling the gradient information generated when updating the model parameters. This method was proposed for computer vision tasks, but was later modified for NLU. We further adjust the algorithm for NMT tasks.

²<https://pytorch.org/>

³<https://github.com/facebookresearch/higher>

Data	Source	Train	Valid	Test
En-Vi	IWSLT’15	133k	768	1268
De-En	IWSLT’14	161k	7.2k	6.7k
Fr-En	IWSLT’16	224k	1080	1133
En-De	WMT’16	4.5m	3.0k	3.0k

Table 1: Dataset source and statistics. Here “k” stands for thousand, and “m” stands for million.

	En-Vi	De-En	Fr-En
Transformer	30.3	34.7	38.2
Adv	31.0	34.8	38.8
FreeAT	31.0	35.2	38.6
FreeLB	31.6	35.3	38.7
SMART	31.5	35.5	38.9
SALT	32.8	36.8	39.7

Table 2: BLEU score on three low-resource datasets. All the baseline results are from our re-implementation. We report the mean of three runs.

◊ *FreeLB* (Zhu et al., 2019) is a “free” large batch adversarial training method. We modify FreeLB to an adversarial regularization method that better fits our need. This algorithm was originally proposed for NLU. We modify the algorithm so that it is also suitable for NMT tasks.

◊ *SMART* (Jiang et al., 2020) is a state-of-the-art fine-tuning method that utilizes smoothness-inducing regularization and Bregman proximal point optimization.

We highlight that we focus on **model generalization** on clean data, instead of adversarial robustness (a model’s ability to defend adversarial attacks). As we will see in the experiments, adversarial training methods (e.g., Adv, FreeAT) suffer from label leakage, and do not generalize as well as adversarial regularization methods.

4.2 Neural Machine Translation

Datasets. We adopt three low-resource datasets and a rich-resource dataset. Dataset statistics are summarized in Table 1. For the low-resource experiments, we use⁴: English-Vietnamese from IWSLT’15, German-English from IWSLT’14, and French-English from IWSLT’16. For the rich-resource experiments, we use the English-German dataset from WMT’16, which contains about 4.5 million training samples.

⁴<https://iwslt.org/>

	RTE Acc	MRPC Acc/F1	CoLA Mcc	SST-2 Acc	STS-B P/S Corr	QNLI Acc	QQP Acc/F1	MNLI-m/mm Acc	Average Score
BERT _{LARGE}	71.1	86.0/89.6	61.8	93.5	89.6/89.3	92.4	91.3/88.4	86.3/86.2	84.0
BERT _{BASE}	63.5	84.1/89.0	54.7	92.9	89.2/88.8	91.1	90.9/88.3	84.5/84.4	81.5
FreeAT	68.0	85.0/89.2	57.5	93.2	89.5/89.0	91.3	91.2/88.5	84.9/85.0	82.6
FreeLB	70.0	86.0/90.0	58.9	93.4	89.7/89.2	91.5	91.4/88.4	85.4/85.5	83.3
SMART	71.2	87.7/91.3	59.1	93.0	90.0/89.4	91.7	91.5/88.5	85.6/86.0	83.8
SALT	72.9	88.4/91.8	61.0	93.6	90.4/90.0	92.0	91.7/88.6	86.1/85.8	84.5

Table 3: Evaluation results on the GLUE development set. All the rows use $BERT_{BASE}$, except the top one, which is included to demonstrate the effectiveness of our model. Best results on each dataset, excluding $BERT_{LARGE}$, are shown in **bold**. Results of $BERT_{BASE}$ (Devlin et al., 2019), $BERT_{LARGE}$ (Devlin et al., 2019), $FreeAT$ (Shafahi et al., 2019), and $FreeLB$ (Zhu et al., 2019) are from our re-implementation. $SMART$ results are from Jiang et al. (2020).

	RTE Acc	MRPC Acc/F1	CoLA Mcc	SST-2 Acc	STS-B P/S Corr	QNLI Acc	QQP Acc/F1	MNLI-m/mm Acc	Average Score
BERT _{BASE}	66.4	84.8/88.9	52.1	93.5	87.1/85.8	90.5	71.2/89.2	84.6/83.4	80.0
FreeLB	70.1	83.5/88.1	54.5	93.6	87.7/86.7	91.8	72.7/89.6	85.7/84.6	81.2
SALT	72.2	85.8/89.7	55.6	94.2	88.0/87.1	92.1	72.8/89.8	85.8/84.8	82.0

Table 4: GLUE test set results on the GLUE evaluation server. All the methods fine-tune a pre-trained $BERT_{BASE}$ model. $FreeAT$ and $SMART$ did not report $BERT_{BASE}$ results in their paper or on the GLUE evaluation server. Model references: $BERT_{BASE}$ (Devlin et al., 2019), $FreeLB$ (Zhu et al., 2019).

	BLEU
Transformer (Vaswani et al., 2017)	28.4
FreeAT (Shafahi et al., 2019)	29.0
FreeLB (Zhu et al., 2019)	29.0
SMART (Jiang et al., 2020)	29.1
SALT	29.6

Table 5: sacreBLEU score on WMT’16 En-De. All the baseline results are from our re-implementation.

Implementation. Recall that to generate adversarial examples, we perturb the word embeddings. In NMT experiments, we perturb both the source-side and the target-side embeddings. This strategy is empirically demonstrated (Sato et al., 2019) to be more effective than perturbing only one side of the inputs. We use *Fairseq*⁵ (Ott et al., 2019) to implement our algorithms. We adopt the Transformer-base (Vaswani et al., 2017) architecture in all the low-resource experiments, except IWSLT’14 De-En. In this dataset, we use a model smaller than Transformer-base by decreasing the hidden dimension size from 2048 to 1024, and decreasing the number of heads from 8 to 4 (while dimension of each head doubles). For the rich-resource experi-

⁵<https://github.com/pytorch/fairseq>

ments, we use the Transformer-big (Vaswani et al., 2017) architecture. Training details are presented in Appendix B.1.

Results. Experimental results for the low-resource experiments are summarized in Table 2. Notice that SMART, which utilizes conventional adversarial regularization, consistently outperforms standard adversarial training (Adv). Similar observations were also reported in Miyato et al. (2017); Sato et al. (2019). This is because Adv generates perturbations using the correct examples, thus, the label information are “leaked” (Kurakin et al., 2017). Additionally, we can see that SALT is particularly effective in this low-resource setting, where it outperforms all the baselines by large margins. In comparison with the vanilla Transformer model, SALT achieves up to 2 BLEU score improvements on all the three datasets.

Table 5 summarizes experiment results on the WMT’16 En-De dataset. We report the sacreBLEU (Post, 2018) score, which is a detokenized version of the BLEU score that better reflects translation quality. We can see that SALT outperforms all the baseline methods by notable margins, and it improves upon the vanilla Transformer model by 1.2 BLEU score.




Figure 1: Relation between BLEU score and different factors on the IWSLT’14 De-En dataset.

4.3 Natural Language Understanding

Datasets. We demonstrate the effectiveness of SALT on the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2019a), which is a collection of nine NLU tasks. The benchmark includes question answering (Rajpurkar et al., 2016), linguistic acceptability (CoLA, Warstadt et al. 2019), sentiment analysis (SST, Socher et al. 2013), text similarity (STS-B, Cer et al. 2017), paraphrase detection (MRPC, Dolan and Brockett 2005), and natural language inference (RTE & MNLI, Dagan et al. 2006; Bar-Haim et al. 2006; Giampiccolo et al. 2007; Bentivogli et al. 2009; Williams et al. 2018) tasks. Dataset details can be found in Table 7 (Appendix B.2).

Implementation. We evaluate our algorithm by fine-tuning a pre-trained BERT-base (Devlin et al., 2019) model. Our implementation is based on the *MT-DNN* code-base (Liu et al., 2019a, 2020b)⁶. Training details are presented in Appendix B.2.

Results. Table 3 summarizes experiment results on the GLUE development set. We can see that SALT outperforms BERT_{BASE} in all the tasks. Further, our method is particularly effective for small datasets, such as RTE, MRPC, and CoLA, where we achieve 9.4, 4.3, and 6.3 absolute improvements, respectively. Comparing with other adversarial training baselines, i.e., FreeAT, FreeLB, and SMART, our method achieves notable improvements in all the tasks.

We highlight that SALT achieves a 84.5 average score, which is significantly higher than that of the vanilla BERT_{BASE} (+3.0) fine-tuning approach. Also, our average score is higher than the scores of baseline adversarial training methods (+1.9, +1.2, +0.7 for FreeAT, FreeLB, SMART, respectively). Moreover, the 84.5 average score is even higher

than fine-tuning BERT_{LARGE} (+0.5), which contains three times more parameters than the backbone of SALT.

Table 4 summarizes results on the GLUE test set. We can see that SALT consistently outperforms BERT_{BASE} and FreeLB across all the tasks.

4.4 Parameter Study

◇ **Robustness to the number of unrolling steps.** From Figure 1a, we can see that SALT is robust to the number of unrolling steps. As such, setting the unrolling steps $K = 1$ or 2 suffices to build models that generalize well.

◇ **Robustness to the perturbation strength.** Unrolling is robust to the perturbation strength within a wide range, as indicated in Figure 1b. Meanwhile, the performance of SMART consistently drops when we increase ϵ from 0.01 to 0.5. This indicates that the unrolling algorithm can withstand stronger perturbations than conventional approaches.

◇ **ℓ_2 constraints vs. ℓ_∞ constraints.** Figure 1c illustrates model performance with respect to different perturbation strength in the ℓ_∞ case. Notice that in comparison with the ℓ_2 case (Figure 1b), SALT achieves the same level of performance, but the behavior of SMART is unstable. Additionally, SALT is stable within a wider range of perturbation strength in the ℓ_2 than in the ℓ_∞ case, which is the reason that we adopt ℓ_2 constraints in the experiments.

We highlight that SALT **does not introduce additional tuning parameter** comparing with conventional adversarial regularization approaches.

4.5 Analysis

◇ **Unrolling reduces bias.** In Figure 3, we visualize the training and the validation error on the STS-B and the SST datasets from the GLUE benchmark. As mentioned, conventional adversarial regulariza-

⁶<https://github.com/microsoft/MT-DNN>




Figure 2: Reliability diagrams on SST. Perfect Calibration: confidence = accuracy; ECE: the lower the better.




Figure 3: Training and validation loss of *SMART* and *SALT* on STS-B (upper) and SST-2 (lower) datasets.

tion suffers from over-strong perturbations, such that the model cannot fit the unperturbed data well. This is supported by the fact that the training loss of *SALT* is smaller than that of *SMART*, which means *SALT* fits the data better. *SALT* also yields a smaller loss than *SMART* on the validation data, indicating that the Stackelberg game-formulated model exhibits better generalization performance.

◇ **Adversarial robustness.** Even though the primary focus of *SALT* is model generalization, we still test its robustness on the Adversarial-NLI (ANLI, Nie et al. 2020) dataset. The dataset contains 163k data, which are collected via a human-and-model-in-the-loop approach. From Table 6, we can see that *SALT* improves model robustness upon conventional methods (i.e., *SMART*).

◇ **Probing experiments.** For each method, we first fine-tune a *BERT*_{BASE} model on the SST-2 dataset. Then, we only tune a prediction head on other

	Dev			
	R1	R2	R3	All
<i>BERT</i> _{BASE}	53.3	43.0	44.7	46.8
<i>SMART</i>	54.1	44.4	45.3	47.8
<i>SALT</i>	56.6	46.2	45.9	49.3
	Test			
	R1	R2	R3	All
<i>BERT</i> _{BASE}	54.1	44.9	46.6	48.4
<i>SMART</i>	54.3	46.4	46.5	48.9
<i>SALT</i>	55.4	47.7	46.7	49.7

Table 6: Experimental results on the ANLI dataset. Model references: *BERT*_{BASE} (Devlin et al., 2019), *SMART* (Jiang et al., 2020).




Figure 4: Probing experiments. Each violin plot is based on 10 runs with different random seeds.

datasets while keeping the representations fixed. Such a method directly measures the quality of representations generated by different models. As illustrated in Fig. 4, *SALT* outperforms the baseline methods by large margins.

◇ **Classification Model Calibration.** Adversarial regularization also helps model calibration (Stutz et al., 2020). A well-calibrated model produces reliable confidence estimation (i.e., confidence \simeq actual accuracy), where the confidence is defined as the maximum output probability calculated by the model. We evaluate the calibration performance of *BERT*_{BASE}, *SMART*, and *SALT* by the Expected Calibration Error (ECE, Niculescu-Mizil and Caruana 2005). We plot the reliability diagram (confidence vs. accuracy) on the SST task in Fig. 2 (see

Appendix C for details). As we can see, BERT_{BASE} and SMART are more likely to make overconfident predictions. SALT reduces ECE, and its corresponding reliability diagram aligns better with the perfect calibration curve.

◇ **Comparison with Unrolled-GAN.** The unrolling technique has been applied to train GANs (Unrolled-GAN, Metz et al. 2017). However, subsequent works find that this approach not necessarily improves training (Grnarova et al., 2018; Tran et al., 2019; Doan et al., 2019). This is because Unrolled-GAN unrolls its discriminator, which has a significant amount of parameters. Consequently, the unrolling algorithm operates on a very large space, rendering the stochastic gradients that are used for updating the discriminator considerably noisy. In SALT, the unrolling space is the sample embedding space, the dimension of which is much smaller than the unrolling space of GANs. Therefore, unrolling is more effective for NLP tasks.

5 Conclusion

We propose SALT, an adversarial regularization method that employs a Stackelberg game formulation. Such a formulation induces a competition between a leader (the model) and a follower (the adversary). In SALT, the leader is in an advantageous position by recognizing the follower’s strategy, and this strategic information is captured by the Stackelberg gradient. We compute the Stackelberg gradient, and hence find the equilibrium of the Stackelberg game, using an unrolled optimization approach. Empirical results NMT and NLU tasks suggest the superiority of SALT to existing adversarial regularization methods.

Broader Impact

This paper proposes Stackelberg Adversarial Regularization (SALT), an adversarial regularized training framework for NLP tasks. Different from Generative Adversarial Networks (GAN), where the target is to attack existing neural network models, or to improve models’ robustness to adversarial attacks, we seek to improve the generalization performance of deep learning models. We demonstrate that the SALT framework can be used for neural machine translation and natural language understanding tasks. In all the experiments, we use publicly available data, and we build our algorithms using public code bases. We do not find any ethical concerns.

References

- Marcin Andrychowicz, Misha Denil, Sergio Gomez Colmenarejo, Matthew W. Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. 2016. [Learning to learn by gradient descent by gradient descent](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 3981–3989.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Roy Bar-Haim, Ido Dagan, Bill Dolan, Lisa Ferro, and Danilo Giampiccolo. 2006. The second PASCAL recognising textual entailment challenge. In *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. 2009. The fifth pascal recognizing textual entailment challenge. In *In Proc Text Analysis Conference (TAC’09)*.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Hao Cheng, Xiaodong Liu, Lis Pereira, Yaoliang Yu, and Jianfeng Gao. 2021. [Posterior differential regularization with f-divergence for improving model robustness](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1078–1089, Online. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Christopher D. Manning, and Quoc Le. 2018. [Semi-supervised sequence modeling with cross-view training](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925, Brussels, Belgium. Association for Computational Linguistics.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. [The pascal recognising textual entailment challenge](#). In *Proceedings of the First International Conference on Machine Learning Challenges: Evaluating Predictive Uncertainty Visual Object Classification, and Recognizing Textual Entailment, MLCW’05*, pages 177–190, Berlin, Heidelberg. Springer-Verlag.
- 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.
- Thang Doan, João Monteiro, Isabela Albuquerque, Bogdan Mazouze, Audrey Durand, Joelle Pineau, and R. Devon Hjelm. 2019. [On-line adaptative curriculum learning for gans](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 3470–3477. AAAI Press.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *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 1126–1135. PMLR.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](#). In *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 1243–1252. PMLR.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague. Association for Computational Linguistics.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Edward Grefenstette, Brandon Amos, Denis Yarats, Phu Mon Htut, Artem Molchanov, Franziska Meier, Douwe Kiela, Kyunghyun Cho, and Soumith Chintala. 2019. Generalized inner loop meta-learning. *arXiv preprint arXiv:1910.01727*.
- Paulina Grnarova, Kfir Y. Levy, Aurélien Lucchi, Thomas Hofmann, and Andreas Krause. 2018. [An online learning approach to generative adversarial networks](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. [On calibration of modern neural networks](#). In *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 1321–1330. PMLR.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- 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, Zhehui Chen, Yuyang Shi, Bo Dai, and Tuo Zhao. 2021. [Learning to defend by learning to attack](#). In *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 577–585. PMLR.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2020. [SMART: Robust and efficient fine-tuning for pre-trained 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.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Lingkai Kong, Haoming Jiang, Yuchen Zhuang, Jie Lyu, Tuo Zhao, and Chao Zhang. 2020. [Calibrated language model fine-tuning for in- and out-of-distribution data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1326–1340, Online. Association for Computational Linguistics.
- Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. 2017. [Adversarial machine learning at scale](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Yan Li, Ethan X Fang, Huan Xu, and Tuo Zhao. 2019. Inductive bias of gradient descent based adversarial training on separable data. *arXiv preprint arXiv:1906.02931*.

- Xiaodong Liu, Hao Cheng, Pengcheng He, Weizhu Chen, Yu Wang, Hoifung Poon, and Jianfeng Gao. 2020a. Adversarial training for large neural language models. *arXiv preprint arXiv:2004.08994*.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019a. [Multi-task deep neural networks for natural language understanding](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496, Florence, Italy. Association for Computational Linguistics.
- Xiaodong Liu, Yu Wang, Jianshu Ji, Hao Cheng, Xueyun Zhu, Emmanuel Awa, Pengcheng He, Weizhu Chen, Hoifung Poon, Guihong Cao, and Jianfeng Gao. 2020b. [The Microsoft toolkit of multi-task deep neural networks for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 118–126, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Dougal Maclaurin, David Duvenaud, and Ryan P. Adams. 2015. [Gradient-based hyperparameter optimization through reversible learning](#). In *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 2113–2122. JMLR.org.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. [Towards deep learning models resistant to adversarial attacks](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. 2017. [Unrolled generative adversarial networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Yifei Min, Lin Chen, and Amin Karbasi. 2020. [The curious case of adversarially robust models: More data can help, double descend, or hurt generalization](#). *arXiv preprint arXiv:2002.11080*.
- Takeru Miyato, Andrew M. Dai, and Ian J. Goodfellow. 2017. [Adversarial training methods for semi-supervised text classification](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Mahdi Pakdaman Naeini, Gregory F. Cooper, and Milos Hauskrecht. 2015. [Obtaining well calibrated probabilities using bayesian binning](#). In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pages 2901–2907. AAAI Press.
- Alexandru Niculescu-Mizil and Rich Caruana. 2005. [Predicting good probabilities with supervised learning](#). In *Machine Learning, Proceedings of the Twenty-Second International Conference (ICML 2005), Bonn, Germany, August 7-11, 2005*, volume 119 of *ACM International Conference Proceeding Series*, pages 625–632. ACM.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. [Scaling neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8024–8035.
- Barak A Pearlmutter and Jeffrey Mark Siskind. 2008. [Reverse-mode ad in a functional framework: Lambda the ultimate backpropagator](#). *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(2):1–36.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). 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

- 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John C. Duchi, and Percy Liang. 2020. [Understanding and mitigating the tradeoff between robustness and accuracy](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 7909–7919. PMLR.
- 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.
- Motoki Sato, Jun Suzuki, and Shun Kiyono. 2019. [Effective adversarial regularization for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 204–210, Florence, Italy. Association for Computational Linguistics.
- Motoki Sato, Jun Suzuki, Hiroyuki Shindo, and Yuji Matsumoto. 2018. [Interpretable adversarial perturbation in input embedding space for text](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4323–4330. ijcai.org.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John P. Dickerson, Christoph Studer, Larry S. Davis, Gavin Taylor, and Tom Goldstein. 2019. [Adversarial training for free!](#) In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3353–3364.
- Qianli Shen, Yan Li, Haoming Jiang, Zhaoran Wang, and Tuo Zhao. 2020. [Deep reinforcement learning with robust and smooth policy](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 8707–8718. PMLR.
- 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, Seattle, Washington, USA. Association for Computational Linguistics.
- David Stutz, Matthias Hein, and Bernt Schiele. 2019. [Disentangling adversarial robustness and generalization](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6976–6987. Computer Vision Foundation / IEEE.
- David Stutz, Matthias Hein, and Bernt Schiele. 2020. [Confidence-calibrated adversarial training: Generalizing to unseen attacks](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 9155–9166. PMLR.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. [Intriguing properties of neural networks](#). In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Ngoc-Trung Tran, Viet-Hung Tran, Ngoc-Bao Nguyen, Linxiao Yang, and Ngai-Man Cheung. 2019. [Self-supervised GAN: analysis and improvement with multi-class minimax game](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13232–13243.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Heinrich Von Stackelberg. 2010. *Market structure and equilibrium*. Springer Science & Business Media.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

- Dilin Wang, ChengYue Gong, and Qiang Liu. 2019b. [Improving neural language modeling via adversarial training](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6555–6565. PMLR.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- 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, New Orleans, Louisiana. Association for Computational Linguistics.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqu Sun, Thomas Goldstein, and Jingjing Liu. 2019. [Freelb: Enhanced adversarial training for language understanding](#). *arXiv preprint arXiv:1909.11764*.

A Virtual Adversarial Training

Virtual adversarial training (VAT, Miyato et al. 2017) solves the following min-max optimization problem:

$$\min_{\theta} \mathcal{F}(\theta, \delta^*) = \mathcal{L}(\theta) + \frac{\alpha}{n} \sum_{i=1}^n \ell_v(x_i, \delta_i^*, \theta),$$

$$\delta_i^* = \operatorname{argmax}_{\|\delta_i\| \leq \epsilon} \ell_v(x_i, \delta_i, \theta),$$

where

$$\ell_v(x_i, \delta_i, \theta) = \text{KL}(f(x_i, \theta) \parallel f(x_i + \delta_i, \theta)).$$

Note that the objective of the minimization problem is a function of both the model parameters and the perturbations.

Because the min problem and the max problem are operating on the same loss function, i.e., the min problem seeks to minimize ℓ_v , while the max problem tries to maximize ℓ_v , this min-max optimization is essentially a zero-sum game. And we can find the game’s equilibrium using gradient descent/ascent algorithms.

Specifically, the adversarial player first generate an initial perturbation δ^0 , and then refines it using K steps of projected gradient ascent, i.e.,

$$\delta^k = \Pi_{\|\cdot\| \leq \epsilon} \left(\delta^{k-1} + \eta \frac{\partial \ell_v(x, \delta^{k-1}, \theta)}{\partial \delta^{k-1}} \right),$$

for $k = 1, \dots, K$.

Here Π denotes projection onto the ℓ_2 -ball or the ℓ_∞ -ball. Empirically, we find that these two choices yield very similar performance, although adversarial training models is robust to ϵ within a wider range when applying the ℓ_2 constraint.

After obtaining the K -step refined perturbation δ^K , we use gradient descent to update the model parameters θ . Concretely, the gradient of the model parameters is computed as

$$\frac{\partial \mathcal{F}(\theta, \delta^K)}{\partial \theta} = \frac{d\ell(f(x_i, \theta), y_i)}{d\theta} + \alpha \frac{\partial \ell_v(x, \delta^K, \theta)}{\partial \theta}. \quad (7)$$

The training algorithm is demonstrated in Algorithm 2.

Note that in this paper, we target for models’ generalization performance on the unperturbed test data, therefore we do not want a strong adversary that “traps” the model parameters to a bad local optima. Most of the existing algorithms achieve this goal by carefully tuning the hyper-parameters

ϵ and K , i.e., a small ϵ usually generates weaker adversaries, so does a small K . However, these heuristics do not work well, and at times δ^K is too strong. Consequently, conventional adversarial training results in undesirable underfitting on the clean data.

Algorithm 2: Virtual Adversarial Training.

Input: \mathcal{D} : dataset; T : total number of training iterations; σ^2 : variance of initial perturbations; K : number of inner training iterations; η : step size to update δ ; Optimizer: optimizer to update θ .

Initialize: model parameters θ ;

for $t = 1, \dots, T$ **do**

for $(x, y) \in \mathcal{D}$ **do**

 Initialize $\delta^0 \sim \mathcal{N}(0, \sigma^2 I)$;

for $k = 1, \dots, K$ **do**

$g^k \leftarrow \partial \ell_v(x_i, \delta_i, \theta) / \partial \delta_i$;

$\delta^k \leftarrow \Pi(\delta^{k-1} + \eta g^k)$;

 Compute the gradient g_θ using

 Eq. 7;

$\theta \leftarrow \text{Optimizer}(g_\theta)$;

Output: θ

B Training Details

B.1 Neural Machine Translation

For the rich-resource WMT’16 En-De dataset, we use the pre-processed data from Ott et al. (2018)⁷. For the low-resource datasets, we use byte-pair encoding (Sennrich et al., 2016) with 10,000 merge operations to build the vocabulary for the IWSLT (’14, ’15, ’16) datasets. We follow the scripts in Ott et al. (2019)⁸ for other pre-processing steps.

We use Adam (Kingma and Ba, 2015) as the leader’s (i.e., the upper level problem that solves for model parameters) optimizer, and we set $\beta = (0.9, 0.98)$. The follower’s (i.e., the lower level problem that solves for perturbations) optimizer is chosen from Adam and SGD, where we observe only marginal empirical differences between these two choices. For low-resource translation, we set the batch size to be equivalent to 64k tokens. For example, when running the experiments on 4 GPUs,

⁷https://github.com/pytorch/fairseq/tree/master/examples/scaling_nmt

⁸<https://github.com/pytorch/fairseq/tree/master/examples/translation>

Corpus	Task	#Train	#Dev	#Test	#Label	Metrics
Single-Sentence Classification (GLUE)						
CoLA	Acceptability	8.5k	1k	1k	2	Matthews corr
SST	Sentiment	67k	872	1.8k	2	Accuracy
Pairwise Text Classification (GLUE)						
MNLI	NLI	393k	20k	20k	3	Accuracy
RTE	NLI	2.5k	276	3k	2	Accuracy
QQP	Paraphrase	364k	40k	391k	2	Accuracy/F1
MRPC	Paraphrase	3.7k	408	1.7k	2	Accuracy/F1
QNLI	QA/NLI	108k	5.7k	5.7k	2	Accuracy
Text Similarity (GLUE)						
STS-B	Similarity	7k	1.5k	1.4k	1	Pearson/Spearman corr

Table 7: Summary of the GLUE benchmark.

	Batch	$\text{lr}_{\text{leader}}$	$\text{lr}_{\text{follower}}$	σ	ϵ	K	Beam	Len-Pen
En-Vi (IWSLT’15)	64k	1×10^{-3}	1×10^{-5}	1×10^{-4}	0.1	1	10	1.0
De-En (IWSLT’14)	64k	1×10^{-3}	1×10^{-4}	1×10^{-4}	0.3	1	9	1.5
Fr-En (IWSLT’16)	64k	1×10^{-3}	1×10^{-5}	1×10^{-5}	0.3	1	10	2.0
En-De (WMT’16)	450k	1×10^{-3}	1×10^{-4}	1×10^{-4}	0.3	1	4	0.6

Table 8: Hyper-parameters for machine translation. Here, σ is the standard deviation of the initial perturbations, ϵ is the perturbation strength, K is the number of unrolling steps, $Beam$ is the size of beam search, and $Len-Pen$ is the length penalty parameter during beam search.

we set the tokens-per-GPU to be 8,000, and we accumulate gradients for 2 steps. For rich-resource translation, we set the batch size to be equivalent to 450k tokens. In all the experiments, we constrain each perturbation according to its sentence-level ℓ_2 norm, i.e., $\|\delta\|_2 \leq \epsilon$. Other hyper-parameters are specified in Table 8.

B.2 Natural Language Understanding

Details of the GLUE benchmark, including tasks, statistics, and evaluation metrics, are summarized in Table 7.

We use Adam as both the leader’s and the follower’s optimizer, and we set $\beta = (0.9, 0.98)$. The learning rate of the leader $\text{lr}_{\text{leader}}$ is chosen from $\{5 \times 10^{-5}, 1 \times 10^{-4}, 5 \times 10^{-4}\}$, and the follower’s learning rate is chosen from $\{1 \times 10^{-5}, \text{lr}_{\text{leader}}\}$. We choose the batch size from $\{4, 8, 16, 32\}$, and we train for a maximum 6 epochs with early-stopping based on the results on the development set. We apply a gradient norm clipping of 1.0. We set the dropout rate in task specific layers to 0.1. We choose standard deviation of initial perturbations σ from $\{1 \times 10^{-5}, 1 \times 10^{-4}\}$, and ℓ_2 constraints

with perturbation strength $\epsilon = 1.0$ are applied. We set the unrolling steps $K = 2$. We report the best performance on each dataset individually.

C Model Calibration

Many applications require trustworthy predictions that need to be not only accurate but also well calibrated (Kong et al., 2020). A well-calibrated model is expected to output prediction confidence comparable to its classification accuracy. For example, given 100 data points with their prediction confidence 0.6, we expect 60 of them to be correctly classified. More precisely, for a data point X , we denote by $Y(X)$ the ground truth label, $\hat{Y}(X)$ the label predicted by the model, and $\hat{P}(X)$ the output probability associated with the predicted label. The calibration error of the predictive model for a given confidence $p \in (0, 1)$ is defined as:

$$\mathcal{E}_p = \left| \mathbb{P} \left[\hat{Y}(X) = Y(X) \mid \hat{P}(X) = p \right] - p \right|. \quad (8)$$

Since Eq. 8 involves population quantities, we usually adopt empirical approximations (Guo et al., 2017) to estimate the calibration error. Specifically,

we partition all data points into 10 bins of equal size according to their prediction confidence. Let \mathcal{B}_m denote the bin with prediction confidence bounded between ℓ_m and u_m . Then, for any $p \in [\ell_m, u_m)$, we define the empirical calibration error as:

$$\widehat{\mathcal{E}}_p = \widehat{\mathcal{E}}_m = \frac{1}{|\mathcal{B}_m|} \left| \sum_{i \in \mathcal{B}_m} [\mathbf{1}(\widehat{y}_i = y_i) - \widehat{p}_i] \right|, \quad (9)$$

where y_i , \widehat{y}_i and \widehat{p}_i are the true label, predicted label and confidence for sample i .

Reliability Diagram is a bar plot that compares $\widehat{\mathcal{E}}_p$ against each bin, i.e., p . A perfectly calibrated would have $\widehat{\mathcal{E}}_p = (\ell_m + u_m)/2$ for each bin.

Expected Calibration Error (ECE) is the weighted average of the calibration errors of all bins (Naeini et al., 2015) defined as:

$$\text{ECE} = \sum_{m=1}^M \frac{|\mathcal{B}_m|}{n} \widehat{\mathcal{E}}_m, \quad (10)$$

where n is the sample size.

We remark that the goal of calibration is to minimize the calibration error without significantly sacrificing prediction accuracy. Otherwise, a random guess classifier can achieve zero calibration error.