

Noise Stability Regularization for Improving BERT Fine-tuning

Hang Hua^{1*†}, Xingjian Li^{2,3*}, Dejing Dou^{2✉}, Chengzhong Xu³, Jiebo Luo¹

¹University of Rochester, Rochester, NY, USA
{hhua2, jluo}@cs.rochester.edu

²Big Data Lab, Baidu Research, Beijing, China
{lixingjian, doudejing}@baidu.com

³Department of Computer Science, University of Macau, Macau, China
czxu@um.edu.mo

Abstract

Fine-tuning pre-trained language models such as BERT has become a common practice dominating leaderboards across various NLP tasks. Despite its recent success and wide adoption, this process is unstable when there are only a small number of training samples available. The brittleness of this process is often reflected by the sensitivity to random seeds. In this paper, we propose to tackle this problem based on the noise stability property of deep nets, which is investigated in recent literature (Arora et al., 2018; Sanyal et al., 2020). Specifically, we introduce a novel and effective regularization method to improve fine-tuning on NLP tasks, referred to as **Layer-wise Noise Stability Regularization (LNSR)**. We extend the theories about adding noise to the input and prove that our method gives a stabler regularization effect. We provide supportive evidence by experimentally confirming that well-performing models show a low sensitivity to noise and fine-tuning with LNSR exhibits clearly higher generalizability and stability. Furthermore, our method also demonstrates advantages over other state-of-the-art algorithms including L²-SP (Li et al., 2018), Mixout (Lee et al., 2020) and SMART (Jiang et al., 2020).

1 Introduction

Large-scale pre-trained language models such as BERT (Devlin et al., 2019) have been widely used in natural language processing tasks (Guu et al., 2020; Liu, 2019; Wadden et al., 2019; Zhu et al., 2020b). A typical process of training a supervised downstream dataset is to fine-tune a pre-trained model for a few epochs. In this process, most of the model’s parameters are reused, while a random initialized task-specific layer is added to adapt the model to the new task.

Fine-tuning BERT has significantly boosted the state of the art performance on natural lan-

guage understanding (NLU) benchmarks such as GLUE (Wang et al., 2018) and SuperGLUE (Wang et al., 2019). However, despite the impressive empirical results, this process remains unstable due to the randomness involved by data shuffling and the initialization of the task-specific layer. The observed instability in fine-tuning BERT was first discovered by Devlin et al. (2019); Dodge et al. (2020), and several approaches have been proposed to solve this problem (Lee et al., 2020; Zhang et al., 2020; Mosbach et al., 2020).

In this study, we consider the fine-tuning stability of BERT from the perspective of the sensitivity to input perturbation. This is motivated by Arora et al. (2018) and Sanyal et al. (2020) who show that noise injected at the lower layers has very little effect on the higher layers for neural networks with good generalizability. However, for a well pre-trained BERT, we find that the higher layers are still very sensitive to the lower layer’s perturbation (as shown in Figure 1), implying that the high level representations of the pre-trained BERT may not generalize well on downstream tasks and consequently lead to instability. This phenomenon coincides with the observation that transferring the top pre-trained layers of BERT slows down learning and hurts performance (Zhang et al., 2020). In addition, Yosinski et al. (2014) also point out that in transfer learning models for object recognition, the lower pre-trained layers learn more general features while the higher layers closer to the output specialize more to the pre-training tasks. We argue that this result also applies to BERT. Intuitively, if a trained model is insensitive to the perturbation of the lower layers’ output, then the model is confident about the output, and vice versa. Based on the above theoretical and empirical results, we propose a simple and effective regularization method to reduce the noise sensitivity of BERT and thus improve the stability and performance of fine-tuned BERT.

†Contribution during internship at Baidu Research.

* Equal contributions. ✉ Correspondence.

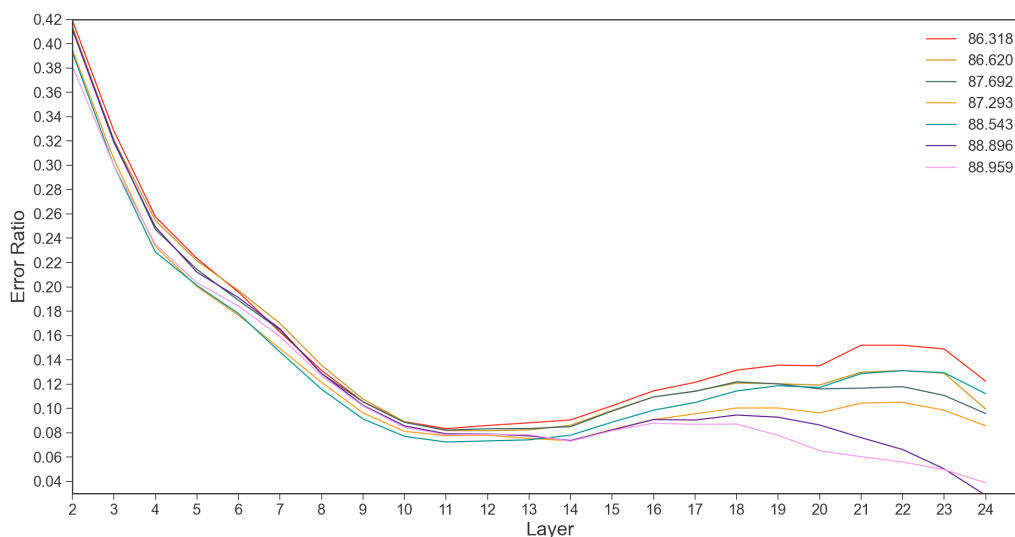


Figure 1: Attenuation of injected noise on the BERT-Large-Uncased model on the MRPC task (X-axis: the layer index. Y-axis: the L^2 norm between the original output and noise perturbed output). A curve starts at the layer where a scaled Gaussian noise is injected to its output whose l_2 norm is set to 5% of the norm of its original output. As it propagates up, the injected noise has a rapidly decreasing effect on the lower layers but becomes volatile on the higher layers, which indicates the poor generalizability and brittleness of the BERT top layers. Moreover, models with higher accuracies (marked in the upper right) usually have lower error ratios or higher noise stability in top layers.

To verify our approach, we conduct extensive experiments on different few-sample (fewer than 10k training samples) NLP tasks, including CoLA (Warstadt et al., 2019), MRPC (Dolan and Brockett, 2005), RTE (Wang et al., 2018; Dagan et al., 2005; Bar-Haim et al., 2006; Giampiccolo et al., 2007), and STS-B (Cer et al., 2017). With the layer-wise noise stability regularization, we obtain strong empirical performance. Compared with other state-of-the-art models, our approach not only improves the fine-tuning stability (with a smaller standard deviation) but also consistently improve the overall performance (with a larger mean, median and maximum).

In summary, our main contributions are:

- We propose a lightweight and effective regularization method, referred to as Layer-wise Noise Stability Regularization (LNSR) to improve the local Lipschitz continuity of each BERT layer and thus ensure the smoothness of the whole model. The empirical results show that the fine-tuned BERT models regularized with LNSR obtain significantly more accurate and stable results. LNSR also outperforms other state-of-the-art methods aiming at stabilizing fine-tuning such as L^2 -SP (Li et al., 2018), Mixout (Lee et al., 2020) and SMART (Jiang et al., 2020).

- We are the first to study the effect of noise stability in NLP tasks. We extend classic theories of adding noise to explicitly constraining the output consistency when adding noise to the input. We theoretically prove that our proposed layer-wise noise stability regularizer is equivalent to a special case of the Tikhonov regularizer, which serves as a stabler regularizer than simply adding noise to the input (Rifai et al., 2011).
- We investigate the relation of the noise stability property to the generalizability of BERT. We find that in general, models with good generalizability tend to be insensitive to noise perturbation; the lower layers of BERT show a better error resilience property but the higher layers of BERT remain sensitive to the lower layers' perturbation (as is depicted in Figure 1).

2 Related Work

2.1 Pre-training

Pre-training has been well studied in machine learning and natural language processing (Erhan et al., 2009, 2010). Mikolov et al. (2013) and Pennington et al. (2014) proposed to use distributional representations (i.e., word embeddings) for indi-

vidual words. Dai and Le (2015) proposed to train a language model or an auto-encoder with unlabeled data and then leveraged the obtained model to finetune downstream tasks. Recently, pre-trained language models, like ELMo (Peters et al., 2018), GPT/GPT-2 (Radford, 2018; Radford et al., 2019), BERT (Devlin et al., 2019), cross-lingual language model (briefly, XLM) (Lample and Conneau, 2019), XLNet (Yang et al., 2019), RoBERTa (Liu et al., 2019) and ALBERT (Lan et al., 2020) have attracted more and more attention in natural language processing communities. The models are first pre-trained on large amount of unlabeled data to capture rich representations of the input, and then applied to the downstream tasks by either providing context-aware embeddings of an input sequence (Peters et al., 2018), or initializing the parameters of the downstream model (Devlin et al., 2019) for fine-tuning. Such pre-training approaches deliver decent performance on natural language understanding tasks.

2.2 Instability in Fine-tuning

Fine-tuning instability of BERT has been reported in various previous works. Devlin et al. (2019) report instabilities when fine-tuning BERT on small datasets and resort to performing multiple restarts of fine-tuning and selecting the model that performs best on the development set. Dodge et al. (2020) performs a large-scale empirical investigation of the fine-tuning instability of BERT. They found dramatic variations in fine-tuning accuracy across multiple restarts and argue how it might be related to the choice of random seed and the dataset size. Lee et al. (2020) propose a new regularization method named Mixout to improve the stability and performance of fine-tuning BERT. Zhang et al. (2020) evaluate the importance of debiasing step empirically by fine-tuning BERT with both BERTAdam and standard Adam optimizer (Kingma and Ba, 2015) and propose a re-initialization method to get a better initialization point for fine-tuning optimization. Mosbach et al. (2020) analyses the cause of fine-tuning instability and propose a simple but strong baseline (small learning rate combined with bias correction).

2.3 Regularization

There has been several regularization approaches to stabilizing the performance of models. Loshchilov and Hutter (2019) propose a decoupled weight decay regularizer integrated in Adam (Kingma and

Ba, 2015) optimizer to prevent neural networks from being too complicate. Gunel et al. (2020) use contrastive learning method to augment training set to improve the generalization performance. In addition, spectral norm (Yoshida and Miyato, 2017; Roth et al., 2019) serves as a general method can also be used to constrain the Lipschitz continuous of matrix, which can increase the stability of generalized neural networks.

There are also several noise-based methods have been proposed to improve the generalizability of pre-trained language models, including SMART (Jiang et al., 2020), FreeLB (Zhu et al., 2020a) and R3F (Aghajanyan et al., 2020). They achieves state of the art performance on GLUE, SNLI (Bowman et al., 2015), SciTail (Khot et al., 2018), and ANLI (Nie et al., 2020) NLU benchmarks. Most of these algorithms employ adversarial training method to improve the robustness of language model fine-tuning. SMART uses an adversarial methodology to encourage models to be smooth within a neighborhoods of all the inputs; FreeLB optimizes a direct adversarial loss through iterative gradient ascent steps; R3F removes the adversarial nature of SMART and optimize the smoothness of the whole model directly. Different from these methods, our proposed method does not adopt the adversarial training strategy, we optimize the smoothness of each layer of BERT directly and thus improve the stability of whole model.

3 Using Noise Stability as a Regularizer

One of the central issues in neural network training is to determine the optimal degree of complexity for the model. A model which is too limited will not sufficiently capture the structure in the data, while one which is too complex will model the noise on the data (the phenomenon of over-fitting). In either case, the performance on new data, that is the ability of the network to generalize, will be poor. The problem can be regarded as one of finding the optimal trade-off between the high bias of a model which is too inflexible and the high variance of a model with too much freedom (Geman et al., 1992; Bishop, 1995; Novak et al., 2018; Bishop, 1991). To control the trade-off of bias against variance of BERT models, we impose an explicit noise regularization method.

3.1 Introduction of Our Method

Denoting the training set as D , we give the general form of optimization objective for a BERT model $f(\cdot; \theta)$ with L layers, as following:

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{(x,y) \sim D} [\mathcal{L}(f(x; \theta), y) + \hat{\mathcal{R}}(\theta)]. \quad (1)$$

To represent $\hat{\mathcal{R}}(\theta)$, we first define the injection position as the *input* of layer b which is denoted as x^b . If the regularization is operated at the *output* of layer r , we can further denote the function between layer b and r as $f^{b,r}$, satisfying that $1 \leq b \leq r \leq L$. To implement the noise stability regularization, we inject a Gaussian-like noise vector ε to x^b and get a neighborhood $x^b + \varepsilon$. Specifically, each element ε_i is independently randomly sampled from a Gaussian distribution with the mean of zero and the standard deviation of σ as $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$. The probability density function of the noise distribution can be written as $p(\varepsilon_i) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{\varepsilon_i^2}{2\sigma^2}}$. Our goal is to minimize the discrepancy between their outputs over $f^{b,r}$ defined as $\|f^{b,r}(x^b + \varepsilon) - f^{b,r}(x^b)\|^2$. In our framework, we use a fixed position b as the position of noise injection and constrain the output distance on all layers following layer b . Denoting the regularization weight corresponding to each $f^{b,r}$ as $\lambda^{b,r}$, given a sample $(x, y) \sim D$, the regularization term is represented by the following formulas:

$$\hat{\mathcal{R}}(\theta) = \sum_{r=b}^L \lambda^{b,r} \|f^{b,r}(x^b + \varepsilon) - f^{b,r}(x^b)\|^2. \quad (2)$$

An overall algorithm is represented in Algorithm 1.

3.2 Theoretical Analysis

Regularization is a kind of commonly used techniques to reduce the function complexity and, as a result, to make the learned model generalize well on unseen examples. In this part, we theoretically prove that the proposed LNSR algorithm has the effects of encouraging the local Lipschitz continuity and imposing a Tikhonov regularizer under different assumptions. For simplicity, we omit the notations about the layer number in this part, denoting f as the target function and x as the input of f parameterized by θ . Given a sample $(x, y) \sim D$, we discuss the general form of the noise stability defined as following:

$$\mathcal{R}(\theta) = E_{\varepsilon} \{ \|f(x + \varepsilon) - f(x)\|^2 \}. \quad (3)$$

Lipschitz continuity. The Lipschitz property reflects the degree of smoothness for a function. Recent theoretical studies on deep learning has revealed the close connection between Lipschitz property and generalization (Bartlett et al., 2017; Neyshabur et al., 2017).

Given a sampled ε , minimizing $\|f(x + \varepsilon) - f(x)\|^2$ is equivalent to minimizing:

$$\frac{\|f(x + \varepsilon) - f(x)\|^2}{\|x + \varepsilon - x\|^2}. \quad (4)$$

Thus the noise stability regularization can be regarded as minimizing the Lipschitz constant in a local region around the input x .

Tikhonov regularizer. The Tikhonov regularizer (Willoughby, 1979) involves constraints on the derivatives of the objective function with respect to different orders. For the simplest first-order case, it can be regarded as imposing robustness and shaping a flatter loss surface at input, which makes the learned function smoother.

Assuming that the magnitude of ε is small, we can expand the first term as a Taylor approximation as:

$$f(x + \varepsilon) = f(x) + \varepsilon \cdot J_f(x) + \frac{1}{2} \varepsilon^T \cdot H_f(x) \cdot \varepsilon + O(\varepsilon^3), \quad (5)$$

where $J_f(x)$ and $H_f(x)$ refer to the Jacobian and Hessian of f with respect to the input x respectively.

Ignoring the higher order term $O(\varepsilon^3)$ and denoting f_k as the k -th output of the function f , we can rewrite the regularizer by substituting Eq. 5 in Eq. 3 as:

$$\begin{aligned} \mathcal{R}(\theta) &= E_{\varepsilon} \{ \|\varepsilon \cdot J_f(x) + \frac{1}{2} \varepsilon^T \cdot H_f(x) \cdot \varepsilon\|^2 \} \\ &= \int \|\varepsilon \cdot J_f(x) + \frac{1}{2} \varepsilon^T \cdot H_f(x) \cdot \varepsilon\|^2 p(\varepsilon) d\varepsilon \\ &= \sum_k \int \|\varepsilon \cdot J_{f_k}(x) + \frac{1}{2} \varepsilon^T \cdot H_{f_k}(x) \cdot \varepsilon\|^2 p(\varepsilon) d\varepsilon. \end{aligned} \quad (6)$$

We define the input vector x as $(x_1, x_2, \dots, x_{d_1})$ and noise vector ε as $\varepsilon = (\varepsilon_1, \varepsilon_2, \dots, \varepsilon_{d_1})$. Assuming that distributions of the noise and the input are irrelevant, and the derivative of f with respect to different elements of the input vector is independent with each other, we expand the second order term

corresponding to the Jacobian as:

$$\begin{aligned}
\Omega_J(f_k) &= \int \|\varepsilon \cdot J_{f_k}(x)\|^2 p(\varepsilon) d\varepsilon \\
&= \int \sum_i (\varepsilon_i \frac{\partial f_k}{\partial x_i})^2 p(\varepsilon) d\varepsilon \\
&= \sum_i (\frac{\partial f_k}{\partial x_i})^2 \int \varepsilon_i^2 p(\varepsilon) d\varepsilon \\
&= \sigma^2 \|J_{f_k}(x)\|^2.
\end{aligned} \tag{7}$$

According to the characteristics of the Gaussian distribution, we also have

$$\int \varepsilon_i \varepsilon_j p(\varepsilon) d\varepsilon = 0 \quad \text{for any } i \neq j. \tag{8}$$

Thus, we can rewrite the second order term corresponding to the Hessian in Eq. 6 as:

$$\begin{aligned}
\Omega_H(f_k) &= \int \|\frac{1}{2} \varepsilon^T \cdot H_{f_k}(x) \cdot \varepsilon\|^2 p(\varepsilon) d\varepsilon \\
&= \int \frac{1}{4} \sum_i (\varepsilon_i^2 \frac{\partial^2 f_k}{\partial x_i^2})^2 p(\varepsilon) d\varepsilon \\
&= \frac{1}{4} \sum_i (\frac{\partial^2 f_k}{\partial x_i^2})^2 \int \varepsilon_i^4 p(\varepsilon) d\varepsilon \\
&= C \|\text{Tr}(H_{f_k}(x)^T H_{f_k}(x))\|^2,
\end{aligned} \tag{9}$$

Where C is a constant independent of the input x . The third term generated from the expansion of Eq. 6 is zero as we have $\int \varepsilon^3 p(\varepsilon) d\varepsilon = 0$. Thus we get

$$\mathcal{R}(\theta) = \sum_k \{\Omega_J(f_k) + \Omega_H(f_k)\}. \tag{10}$$

Considering that the input and output of the function f are both scalar variable, the Tikhonov regularization (Willoughby, 1979) takes the general form as:

$$\mathcal{R}_{\mathcal{T}}(\theta) = \sum_r \int h_r(x) (\frac{\partial^r f}{\partial x^r})^2 dx. \tag{11}$$

Eq. 10 shows that our proposed regularizer ensuring the noise stability is equivalent to a special case of the Tikhonov regularizer, where we involve the first and second order derivatives of the objective function f .

An alternative for improving the robustness is to directly add noise to the input, without explicitly constraining the output stability. (Rifai et al., 2011) has derived that adding noise to the input has the effect of penalizing both the L_2 -norm of

the Jacobian $\|J_f(x)\|^2$ and the trace of the Hessian $\text{Tr}(H_f(x))$, whereas the Hessian term is not constrained to be positive. While the regularizer brought by our proposed LNSR is guaranteed to be positive by involving the sum of squares of the first and second order derivatives. Moreover, our work relaxes the assumption of MSE regression loss required by (Rifai et al., 2011). By imposing the explicit constraint of noise stability on middle layer representations, we extend the theoretical understanding of noise stability into deep learning algorithms.

Algorithm 1 Layer-wise Noise Stability Regularization (LNSR)

Input: Training set D , perturbation bound δ , learning rate τ , number of layers L , number of training epochs N , function f and its corresponding parameters θ , the position of noise injection b , and regularization weights for each layer $\{\lambda^b, \dots, \lambda^L\}$.

- 1: Initialize θ
- 2: **for** epoch=1, 2, ..., N **do**
- 3: **for** minibatch $B \sim D$ **do**
- 4: $\hat{\mathcal{R}} \leftarrow 0$
- 5: **for** each $x \in B$ **do**
- 6: $\varepsilon \sim \mathcal{N}(0, \sigma^2)$
- 7: $\tilde{x} \leftarrow x + \varepsilon$
- 8: forward pass given x and \tilde{x} as inputs
- 9: **for** $r = b, b + 1, \dots, L$ **do**
- 10: $\hat{\mathcal{R}} \leftarrow \hat{\mathcal{R}} + \lambda^{b,r} \|f^{b,r}(x) - f^{b,r}(\tilde{x})\|^2$
- 11: **end for**
- 12: **end for**
- 13: $g \leftarrow \frac{1}{|B|} \sum_{(x,y)} \nabla[\mathcal{L}(f(x; \theta), y) + \hat{\mathcal{R}}]$
- 14: $\theta \leftarrow \theta - \tau g$
- 15: **end for**
- 16: **end for**

Output: θ

4 Experiments

In this section, we experimentally demonstrate the effectiveness of LNSR method on text classification tasks over other regularization methods, and confirm that the insensitivity to noise promotes the generalizability and stability of BERT.

4.1 Data

We conduct experiments on four few-sample (less than 10k training samples) text classification tasks

	RTE			MRPC			CoLA			STS-B		
	mean	std	max	mean	std	max	mean	std	max	mean	std	max
FT (Devlin et al., 2019)	70.13	1.84	72.56	87.57	0.92	89.16	60.54	1.49	62.59	89.38	0.53	90.23
L ² -SP (Li et al., 2018)	70.58	1.29	73.28	87.74	0.86	88.95	60.19	1.42	63.89	89.25	0.62	90.14
Mixout(Lee et al., 2020)	71.35	1.66	74.36	87.63	0.62	88.91	63.12	1.68	65.12	89.58	0.35	90.11
SMART(Jiang et al., 2020)	72.23	2.41	75.45	87.86	0.63	89.09	63.16	1.17	65.21	90.11	0.33	90.83
LNSR(ours)	73.31	1.55	76.17	88.50	0.56	90.02	63.35	1.05	65.99	90.23	0.31	90.97

Table 1: The mean, standard deviation, and maximum performance on the development set of RTE, MRPC, CoLA, and STS-B task across 25 random seeds when fine-tuning the BERT-Large model with various regularization methods. FT refers to the standard BERT fine-tuning. Standard deviation: lower is better.

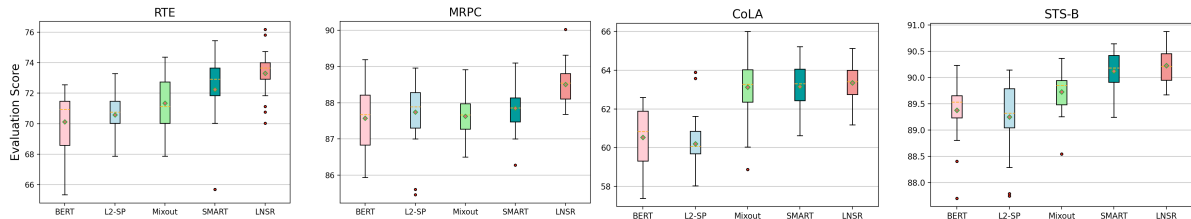


Figure 2: Performance distribution box plot of each model on the four tasks from 25 random seeds.

of GLUE ¹, the datasets are described below and summarized in Appendix A Table 4.

Corpus of Linguistic Acceptability (CoLA (Warstadt et al., 2019)) consists of English acceptability judgments drawn from books and journal articles on linguistic theory. Each example is a sequence of words annotated with whether it is a grammatical English sentence. This is a binary classification task and Matthews correlation coefficient (MCC) (Matthews, 1975) is used to evaluate the performance.

Microsoft Research Paraphrase Corpus (MRPC (Dolan and Brockett, 2005)) is a corpus of sentence pairs with human annotations for whether the sentences in the pair are semantically equivalent. The evaluation metrics is the average of F1 and Accuracy.

Recognizing Textual Entailment (RTE (Wang et al., 2018)) (Dagan et al., 2005) (Bar-Haim et al., 2006) (Giampiccolo et al., 2007) is a corpus of textual entailment, and each example is a sentence pair annotated whether the first entails the second. The evaluation metrics is Accuracy.

Semantic Textual Similarity Benchmark (STS-B (Cer et al., 2017)) is a regression task. Each example is a sentence pair and is human-annotated with a similarity score from 1 to 5; the task is to predict these scores. The evaluation metrics is the average of Pearson and Spearman correlation coef-

ficients.

4.2 Baseline Models

We use BERT (Devlin et al., 2019), a large-scale bidirectional pre-trained language model as the base model in all experiments. We adopt pytorch edition implemented by Wolf et al. (2019).

Fine-tuning. We use the standard BERT fine-tuning method described in Devlin et al. (2019).

L²-SP (Li et al., 2018) is a regularization scheme that explicitly promotes the similarity of the final solution with the initial model. It is usually used for preventing pre-trained models from catastrophic forgetting. We adopt the form of $\Omega(w) = \frac{\alpha}{2} \|w_s - w_s^0\|^2 + \frac{\beta}{2} \|w_{\bar{s}}\|^2$.

Mixout (Lee et al., 2020) is a stochastic regularization technique motivated by Dropout (Srivastava et al., 2014) and DropConnect (Wan et al., 2013). At each training iteration, each model parameter is replaced with its pre-trained value with probability p . The goal is to improve the generalizability of pre-trained language models.

SMART (Jiang et al., 2020) imposes a smoothness regularizer inducing an adversarial manner to control the model complexity at the fine-tuning stage. It also employs a class of Bregman proximal point optimization methods to prevent the model from aggressively updating during fine-tuning.

¹<https://gluebenchmark.com/>

4.3 Experimental Setup

Our model is implemented using Pytorch based on Transformers framework². Specifically, we use the learning setup and hyperparameters recommended by (Devlin et al., 2019). We use Huggingface edition Adam (Kingma and Ba, 2015) optimizer (without bias correction) with learning rate of 2×10^{-5} , $\beta_1 = 0.9$, $\beta_2 = 0.999$, and warmup over the first 10% steps of the total steps. We fine-tune the entire model (340 million parameters), of which the vast majority start as pre-trained weights (BERT-Large-Uncased) and the classification layer (2048 parameters). Weights of the classification layer are initialized with $\mathcal{N}(0, 0.02^2)$. We train with a batch size of 32 for 3 epochs. More details of our experimental setup are described in Appendix A.

4.4 Overall Performance

Table 1 shows the results of all the models on selected GLUE datasets. We train each dataset over 25 random seeds. To implement our LNSR, we uniformly inject noise at the first layer on BERT-large for the comparison with baseline models. As we can see from the table, our model outperforms all the baseline models in mean and max values, which indicates the stronger generalizability of our model against other baseline models. The p-values between the accuracy distributions of standard BERT fine-tuning and our model are calculated to verify whether the improvements are significant. We obtain very small p-values in all tasks: RTE: 9.7×10^{-7} , MRPC: 2.3×10^{-4} , CoLA: 4.7×10^{-8} , STS-2: 3.3×10^{-8} .

Standard deviation is an indicator of the stability of models' performance and higher std means more sensitive to random seeds. Our model shows a lower standard deviation on each task, which means our model is less sensitive to random seeds than other models. Figure 2 presents a clearer illustration. To sum up, our proposed method can effectively improve the performance and stability of fine-tuning BERT.

5 Analysis

5.1 Ablation Study

To verify the effectiveness of our proposed LNSR model, we conduct several ablation experiments including fine-tuning with more training epochs and

²<https://huggingface.co/transformers/index.html>

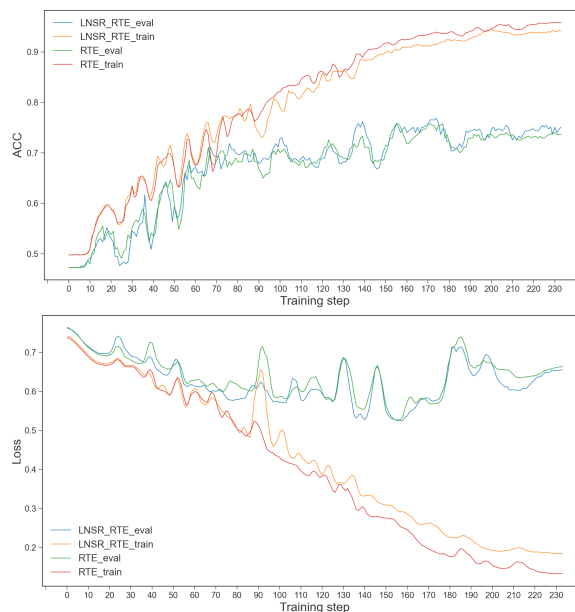


Figure 3: The top-1 accuracy (top) and loss (bottom) curve for fine-tuning BERT with and w/o LNSR.

noise perturbation without regularization (we inject noise directly to the output of a specific layer, and then use the perturbed representation to conduct propagation and then calculate loss, this process is similar to a vector-space represent augmentation). The results are shown in Table 2. We observe that benefit obtained by longer training is limited. Similarly, fine-tuning with noise perturbation only achieves slightly better results on two of these tasks, showing that simply adding noise without an explicit restriction on outputs may not be sufficient to obtain good generalizability. While BERT models with LNSR perform better on each task. This verifies our claim that LNSR can promote the stability of BERT fine-tuning and meanwhile improve the generalizability of the BERT model.

5.2 Effects on the Generalizability of Models

We verify the effects of our proposed method on the generalizability of BERT models in two ways – generalization gap and models' performance on fewer training samples. Due to the limited data and the extremely high complexity of BERT model, bad fine-tuning start point makes the adapted model overfit the training data and does not generalize well to unseen data. Generalizability of models can be intuitively reflected by generalization gap and models' performance on fewer training samples.

Table 3 shows the mean training Acc, mean evaluation Acc and generalization gap of different models on each task. As we can see from the table, fine-tuning with LNSR can effectively narrow the

	RTE			MRPC			CoLA			STS-B		
	mean	std	max	mean	std	max	mean	std	max	mean	std	max
FT	70.13	1.84	72.56	87.57	0.92	89.16	61.56	1.34	64.10	89.38	0.53	90.23
FT (4 Epochs)	70.69	1.97	73.65	88.15	0.65	89.21	60.69	1.24	62.09	89.29	0.56	90.12
FT+Noise	70.62	1.56	72.93	87.95	0.83	89.33	60.18	1.58	62.59	89.34	0.51	90.11
LNSR(ours)	73.31	1.55	76.17	88.50	0.56	90.02	63.35	1.05	65.99	90.23	0.31	90.97

Table 2: Ablation study of LNSR on each task, we report the mean evaluation scores and standard deviation and max value across 25 random seeds. FT refers to standard BERT fine-tuning.

	RTE	MRPC	CoLA	STS-B
	train / eval / gap	train / eval / gap	train / eval / gap	train / eval / gap
FT (Devlin et al., 2019)	95.89/70.13/25.76	96.57/87.57/9.00	97.71/61.56/36.25	98.31/89.38/ 8.93
LNSR(ours)	90.72/ 73.31/17.41	96.68/88.50/8.18	93.44/ 63.35/30.09	98.45/90.23/ 8.22

Table 3: Comparison of the generalizability performance of different models. We report the mean training Acc and evaluation Acc and the generalizability gap (training Acc - evaluation Acc) of each model across 20 random seeds.

generalization gap, and achieve higher evaluation score. The effect of narrowing generalization gap is also reflected in Figure 3 where we can see the higher evaluation accuracy and lower evaluation loss.

We sample subsets from the two relatively larger datasets CoLA (8.5k training samples) and STS-B (7k training samples) with the sampling ratio of 0.15, 0.3 and 0.5. As is shown in Figure 4, fine-tuning with LNSR shows clear advantage on fewer training samples, suggesting LNSR can effectively promote the model’s generalizability.

5.3 Sensitivity to the Position of Noise Injection

We briefly discuss about the sensitivity to the position of noise injection as it is a pre-determined hyperparameter of our method. As is shown in Figure 5 in Appendix A, we observe that the performance of LNSR does not fluctuate much as the position of noise injection changes. All injection positions bring significant improvements over vanilla fine-tuning. Note that, with LNSR, noise injection to the lower layers usually leads to relatively higher accuracy and stability, implying that LNSR may be more effective when it affects both the lower and higher layers of the network.

5.4 Relationship to Previous Noise-based Approaches

Our method is related to SMART (Jiang et al., 2020), FreeLB (Zhu et al., 2020a) and R3F (Agha-

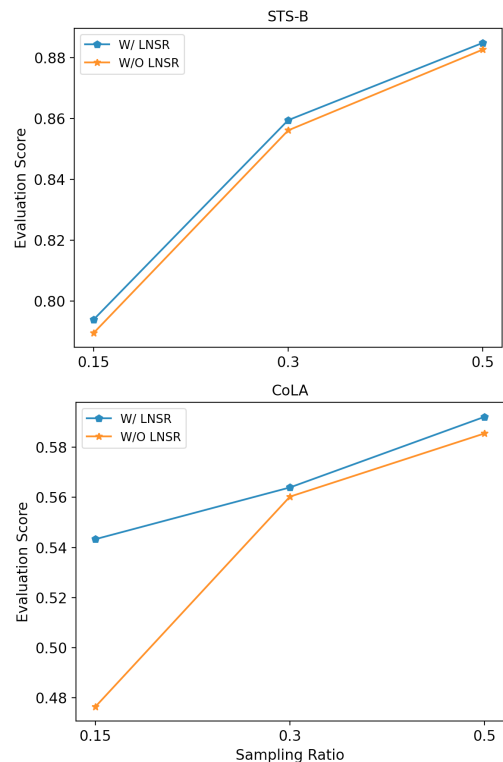


Figure 4: Mean evaluation score comparison of fine-tuning BERT with and w/o LNSR on fewer training samples of the CoLA and STS-B tasks. The mean values are calculated over 20 random seeds.

janyan et al., 2020). As is mentioned before, most of these approaches employ adversarial training strategies to improve the robustness of BERT fine-tuning. SMART solves supremum by using an adversarial methodology to achieve the largest KL di-

vergence with an ϵ -ball, FreeLB optimizes a direct adversarial loss $\mathcal{L}_{FreeLB}(\theta) = \sup_{\Delta\theta: |\Delta\theta| \leq \epsilon} \mathcal{L}(\theta + \Delta\theta)$ through iterative gradient ascent steps, while R3F removes the adversarial nature of SMART and optimize the smoothness of the whole model directly.

Compared with this sort of adversarial based algorithms, our method is easier to implement and provides a relatively rigorous theoretical guarantee. The design of layer-wise regularization is sensible that it exploits the characteristics of hierarchical representations in modern deep neural networks. Studies in knowledge distillation have shown similar experience that imitating through middle layer representations (Adriana et al., 2015; Zagoruyko and Komodakis, 2016) performs better than aligning the final outputs (Hinton et al., 2015). Moreover, LNSR allows us to use different regularization weights for different layers (we use fixed weight 1 on all layers in this paper). We will leave the exploitation in future work.

6 Conclusion

In this paper, we propose the Layer-wise Noise Stability Regularization (LNSR) as a lightweight and effective method to improve the generalizability and stability when fine-tuning BERT on few training samples. Our proposed LNSR method is a general technique that improves model output stability while maintaining or improving the original performance. Furthermore, we provide a theoretically analysis of the relationship of our model to the Lipschitz continuity and Tikhonov regularizer. Extensive empirical results show that our proposed method can effectively improve the generalizability and stability of the BERT model.

7 Acknowledgements

Hang Hua would like to thank Jeffries for supporting his research.

References

Romero Adriana, Ballas Nicolas, K Samira Ebrahimi, Chassang Antoine, Gatta Carlo, and B Yoshua. 2015. Fitnets: Hints for thin deep nets. *Proc. ICLR*.

Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2020. Better fine-tuning by reducing representational collapse. *ArXiv*, abs/2008.03156.

Sanjeev Arora, R. Ge, Behnam Neyshabur, and Yi Zhang. 2018. Stronger generalization bounds

for deep nets via a compression approach. *ArXiv*, abs/1802.05296.

Roy Bar-Haim, I. Dagan, B. Dolan, Lisa Ferro, Danilo Giampiccolo, and B. Magnini. 2006. The second pascal recognising textual entailment challenge.

Peter Bartlett, Dylan J Foster, and Matus Telgarsky. 2017. Spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1706.08498*.

Chris Bishop. 1991. Improving the generalization properties of radial basis function neural networks. *Neural computation*, 3(4):579–588.

Chris M Bishop. 1995. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. *ArXiv*, abs/1508.05326.

Daniel Matthew Cer, Mona T. Diab, Eneko Agirre, I. Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. *ArXiv*, abs/1708.00055.

I. Dagan, Oren Glickman, and B. Magnini. 2005. The pascal recognising textual entailment challenge. In *MLCW*.

Andrew M. Dai and Quoc V. Le. 2015. Semi-supervised sequence learning. In *NIPS*.

J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.

Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah A. Smith. 2020. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *ArXiv*, abs/2002.06305.

W. Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *IWP@IJCNLP*.

D. Erhan, Aaron C. Courville, Yoshua Bengio, and Pascal Vincent. 2010. Why does unsupervised pre-training help deep learning? *J. Mach. Learn. Res.*, 11:625–660.

D. Erhan, Pierre-Antoine Manzagol, Yoshua Bengio, S. Bengio, and P. Vincent. 2009. The difficulty of training deep architectures and the effect of unsupervised pre-training. In *AISTATS*.

Stuart Geman, Elie Bienenstock, and René Doursat. 1992. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58.

- Danilo Giampiccolo, B. Magnini, I. Dagan, and W. Dolan. 2007. The third pascal recognizing textual entailment challenge. In *ACL-PASCAL@ACL*.
- Beliz Gunel, Jingfei Du, Alexis Conneau, and Ves Stoyanov. 2020. Supervised contrastive learning for pre-trained language model fine-tuning. *ArXiv*, abs/2011.01403.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: Retrieval-augmented language model pre-training. *ArXiv*, abs/2002.08909.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- 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 *ACL*.
- Tushar Khot, A. Sabharwal, and Peter Clark. 2018. Scitail: A textual entailment dataset from science question answering. In *AAAI*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *ArXiv*, abs/1901.07291.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations. *ArXiv*, abs/1909.11942.
- Cheolhyoung Lee, Kyunghyun Cho, and Wanmo Kang. 2020. Mixout: Effective regularization to fine-tune large-scale pretrained language models. *ArXiv*, abs/1909.11299.
- Xuhong Li, Y. Grandvalet, and F. Davoine. 2018. Explicit inductive bias for transfer learning with convolutional networks. In *ICML*.
- Y. Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, M. Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Yang Liu. 2019. Fine-tune bert for extractive summarization. *ArXiv*, abs/1903.10318.
- I. Loshchilov and F. Hutter. 2019. Decoupled weight decay regularization. In *ICLR*.
- Brian W Matthews. 1975. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, G. S. Corrado, and J. Dean. 2013. Distributed representations of words and phrases and their compositionality. *ArXiv*, abs/1310.4546.
- Marius Mosbach, Maksym Andriushchenko, and D. Klakow. 2020. On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines. *ArXiv*, abs/2006.04884.
- Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nathan Srebro. 2017. Exploring generalization in deep learning. *arXiv preprint arXiv:1706.08947*.
- 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*. Association for Computational Linguistics.
- Roman Novak, Y. Bahri, D. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. 2018. Sensitivity and generalization in neural networks: an empirical study. *ArXiv*, abs/1802.08760.
- Jeffrey Pennington, R. Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *ArXiv*, abs/1802.05365.
- Jason Phang, Thibault Févry, and Samuel R. Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *ArXiv*, abs/1811.01088.
- A. Radford. 2018. Improving language understanding by generative pre-training.
- A. Radford, Jeffrey Wu, R. Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Salah Rifai, Xavier Glorot, Yoshua Bengio, and Pascal Vincent. 2011. Adding noise to the input of a model trained with a regularized objective. *arXiv preprint arXiv:1104.3250*.
- K. Roth, Yannic Kilcher, and Thomas Hofmann. 2019. Adversarial training generalizes data-dependent spectral norm regularization. *ArXiv*, abs/1906.01527.
- Amartya Sanyal, P. Torr, and Puneet K. Dokania. 2020. Stable rank normalization for improved generalization in neural networks and gans. *ArXiv*, abs/1906.04659.

- Nitish Srivastava, Geoffrey E. Hinton, A. Krizhevsky, Ilya Sutskever, and R. Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15:1929–1958.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, L. Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*.
- David Wadden, Ulme Wennberg, Yi Luan, and Hananeh Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. In *EMNLP/IJCNLP*.
- Li Wan, Matthew D. Zeiler, Sixin Zhang, Y. LeCun, and R. Fergus. 2013. Regularization of neural networks using dropconnect. In *ICML*.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, F. Hill, Omer Levy, and Samuel R. Bowman. 2019. Superglue: A stickier benchmark for general-purpose language understanding systems. *ArXiv*, abs/1905.00537.
- Alex Wang, Amanpreet Singh, Julian Michael, F. Hill, Omer Levy, and Samuel R. Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *ArXiv*, abs/1804.07461.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Ralph A Willoughby. 1979. Solutions of ill-posed problems (an tikhonov and vy arsenin). *SIAM Review*, 21(2):266.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Z. Yang, Zihang Dai, Yiming Yang, J. Carbonell, R. Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*.
- Y. Yoshida and Takeru Miyato. 2017. Spectral norm regularization for improving the generalizability of deep learning. *ArXiv*, abs/1705.10941.
- J. Yosinski, J. Clune, Yoshua Bengio, and Hod Lipson. 2014. How transferable are features in deep neural networks? In *NIPS*.
- Sergey Zagoruyko and Nikos Komodakis. 2016. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*.
- Tianyi Zhang, Felix Wu, Arzoo Katiyar, Kilian Q. Weinberger, and Yoav Artzi. 2020. Revisiting few-sample bert fine-tuning. *ArXiv*, abs/2006.05987.
- C. Zhu, Y. Cheng, Zhe Gan, S. Sun, T. Goldstein, and Jing jing Liu. 2020a. FreeLb: Enhanced adversarial training for natural language understanding. *arXiv: Computation and Language*.
- Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, W. Zhou, H. Li, and T. Liu. 2020b. Incorporating bert into neural machine translation. *ArXiv*, abs/2002.06823.

A Experimental Details

The model we use for experiments in section 4 is the standard BERT large model with 24 layers stacked Transformers (Vaswani et al., 2017) encoder, 1024 hidden size, and 16 self-attention heads. We initialize the pre-trained part of the model with BERT-Large-Uncased-Whole-Word-Masking weight. The final layer is a classification layer with 2048 parameters which contains 0.0006% of the total number of parameters in the model. We initialize the last layer with $\mathcal{N}(0, 0.02^2)$ and each bias is 0. For the position of noise injection, we uniformly chose the first layer as the noise regularization start point. In the sensitivity to the position of noise injection analysis section, we also try injecting noise from the different layers as is shown in Figure 5. As for the baseline model Mixout, we use the code from the Github repository <https://github.com/bloodwass/mixout.git>. The other baseline models are implemented by ourselves.

Table 4 summarizes dataset statistics used in this work. We use the standard GLUE benchmark datasets downloaded from <https://gluebenchmark.com/tasks>.

B Other Experimental Reports

We also report the maximum value we get during fine-tuning BERT with our proposed LNSR regularizer among a large number of random seeds and several noise injection position, since the maximum value can also reflect the ability of the learning algorithm to reach an optimal point. The results are shown in Table 5, and we can see that on some tasks, fine-tuning BERT with LNSR is even competitive with fine-tuning state-of-the-art models which adopt more powerful modern architectures and pre-training strategies.

	RTE	MRPC	CoLA	STS-B
Task	NLI	Paraphrase	Acceptability	Similarity
Metrics	Accuracy	$\frac{\text{Accuracy}+\text{F1}}{2}$	Matthews corr	Pearson/Spearman corr
# of labels	2	2	2	1
# of training samples	2.5k	3.7k	8.6k	7k
# of validation samples	276	408	1k	1.5k
# of test samples	3k	1.7k	1k	1.4k

Table 4: The summarization of the datasets used in this work.

	RTE	MRPC	CoLA	STS-B
BERT (Devlin et al., 2019)	70.4	88.0	60.6	90.0
BERT (Phang et al., 2018)	70.0	90.7	62.1	90.9
LNSR (ours)	79.1	90.4	68.1	91.0
XLNet (Yang et al., 2019)	83.8	89.2	63.6	91.8
RoBERTa (Liu et al., 2019)	86.6	90.9	68.0	92.4
ALBERT (Lan et al., 2020)	89.2	90.9	71.4	93.0

Table 5: We report the maximum value we get when fine-tuning the LNSR model on different noise injection position and random seeds on the four tasks. On some tasks, BERT (standard BERT-large-uncased (Devlin et al., 2019)) with LNSR even become competitive with some newly proposed powerful models (bottom rows)

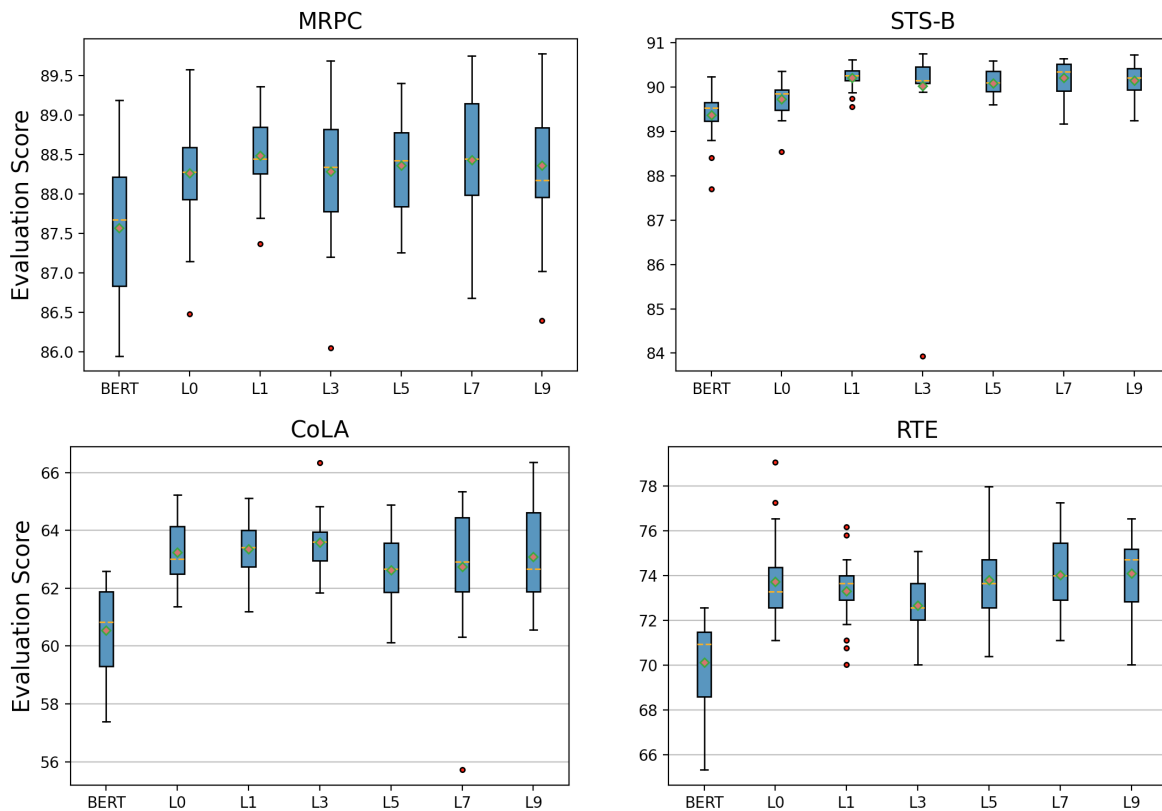


Figure 5: Performance distribution box plot of each model on the four tasks across 25 random seeds.