

Layer-wise Regularized Dropout for Neural Language Models

Shiwen Ni¹, Min Yang^{1,2†}, Ruifeng Xu³, Chengming Li⁴, Xiping Hu⁴

¹Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences

²Key Laboratory of Intelligent Education Technology and Application of Zhejiang Province,
Zhejiang Normal University

³Harbin Institute of Technology, Shenzhen ⁴Shenzhen MSU-BIT University

{sw.ni, min.yang}@siat.ac.cn; xuruifeng@hit.edu.cn; {licm, huxp}@smbu.edu.cn

Abstract

Among the various pre-trained neural language models that are popular today, dropout is already an indispensable regularization technique. To solve the inconsistency between training and inference caused by the randomness of dropout, some studies use consistency training to regularize dropout at the output layer. In this paper, we propose a novel **Layer-wise Regularized Dropout (LR-Drop)**, which is specially designed for Transformer-based Language models. Specifically, LR-Drop layer-wise regularizes each Transformer layer using the consistency training strategy. Each training sample passes through the two siamese sub-models sampled by dropout, and then LR-Drop forces the hidden states, multi-head attention matrices, and output distribution of the two siamese sub-models to be consistent. The proposed LR-Drop can be regarded as a “self-distillation” framework, in which each sub-model generated by dropout is the other’s “teacher” model and “student” model. Through extensive experiments on 8 natural language understanding datasets, 6 neural machine translation datasets, and 1 abstractive summarization dataset (a total of 15 datasets), we show that LR-Drop achieves superior performances, including state-of-the-art results.

Keywords: regularization, dropout, self-distillation, transformer

1. Introduction

In recent years, pre-trained language models (PLMs) based on the Transformer architecture have revolutionized the field of natural language processing (NLP) by achieving state-of-the-art performance on a wide range of NLP tasks. These models, such as BERT (Bidirectional Encoder Representations from Transformers) (Kenton and Toutanova, 2019), ALBERT (A Lite BERT) (Lan et al., 2019), XLNet (Yang et al., 2019), RoBERTa (Liu et al., 2019), and ELECTRA (Efficiently Learning an Encoder that Classifies Token Replacements Accurately) (Clark et al., 2019b), have demonstrated their effectiveness in tasks such as text classification, named entity recognition, sentiment analysis, machine translation, question answering, and more.

One of the key reasons for the success of these Transformer-based PLMs is their ability to capture contextualized representations of words and sentences. By leveraging the self-attention mechanism, these models can efficiently encode the relationships between different words in a sentence, allowing them to capture long-range dependencies and context. The pre-training stage involves training the models on large amounts of unlabeled text, followed by fine-tuning on specific downstream tasks using labeled data. This transfer learning approach has proven to be highly effective, as the pre-trained models can leverage the knowledge learned from the vast amount of unlabeled data to perform well on a variety of NLP tasks.

To prevent overfitting and improve the generalization ability of PLMs, dropout regularization techniques (Srivastava et al., 2014) are commonly employed during both the pre-training and fine-tuning stages. Dropout randomly deactivates a portion of the neural units during training, effectively creating an ensemble of sub-models. This ensemble approach helps in reducing over-reliance on specific units and encourages the model to learn more robust and generalizable representations. However, the use of dropout introduces a challenge in terms of inconsistency between training and inference. During training, dropout is applied to create the ensemble, but during inference, the full model without dropout is used, leading to a mismatch in behavior.

Several studies (Ma et al., 2016; Zolna et al., 2018) have highlighted this inconsistency and its potential impact on model performance. They have proposed methods to address this issue by introducing L2 regularization to the hidden unit state. However, the effectiveness of this approach is limited, and it does not fully resolve the inconsistency problem. To tackle this challenge more effectively, recent research (Wu et al., 2021) has introduced a novel consistency training method called R-Drop. R-Drop aims to align the output distributions of identical data samples processed by different sub-models created through dropout. It involves performing two forward passes for each data sample, with each pass handled by a distinct sub-model that randomly deactivates some hidden units. By minimizing the bidirectional Kullback-Leibler (KL) divergence between the output distributions of these

[†]Corresponding author

two sub-models, R-Drop encourages consistency in the predictions made by the ensemble. This approach provides a more robust and consistent regularization of dropout, addressing the inconsistency issue between training and inference.

In addition to regulating dropout at the output layer, it is also important to ensure consistency in other representations within the PLM. For instance, the multi-head attention mechanism, which is a crucial component of Transformer-based models, typically employs dropout. Previous studies (Clark et al., 2019a) have shown that the attention weight matrix captures substantial linguistic knowledge. Therefore, it is essential to maintain consistency between the multi-head attention matrices of different sub-models to preserve the learned linguistic knowledge. By extending the principles of R-Drop, we proposed LR-Drop to introduce regular constraints into each Transformer layer of the model.

In particular, we formulate three loss functions to regulate different representations from PLMs layers: 1) the hidden states and 2) multi-head attention matrices extracted from the Transformer layer; 3) the output distributions generated by the prediction layer. The multi-head attention in PLMs typically employs dropout, and previous studies (Clark et al., 2019a) have demonstrated that the attention weight matrix can acquire a substantial amount of linguistic knowledge, hence we ensure the consistency between the two multi-head attentions.

To summarize, the main contributions of this paper are as follows:

- In this work, we propose the layer-wise regularized dropout (LR-Drop), a simple but effective regularization technique built upon dropout, designed for Transformer-based pre-trained language models.
- For the special structure of Transformer-based pre-trained language models, we are the first to propose Transformer-layer regularization, which includes regularization for hidden states and multi-headed attention.
- Our LR-Drop does not introduce additional model parameters and does not change the original architecture of the language model.
- By conducting rigorous experiments on 8 natural language understanding datasets, 6 neural machine translation datasets, and 1 abstractive summarization dataset, we provide evidence that LR-Drop excels in performance, even achieving state-of-the-art results.

2. Related Work

2.1. Regularization Methods

The susceptibility of large and deep neural network models to overfitting is a well-established fact. It has been observed that the most effective models are typically large ones, but they are also paired with appropriate regularization techniques. A plethora of regularization techniques have been suggested to enhance the generalization capacity of these models. (Krogh and Hertz, 1992) introduced simple weight decay as a regularization technique to improve generalizability. (Kang et al., 2016) proposed the Shakeout method, which randomly enhances or inverts the contribution of each cell to the next layer, effectively applying L1 and L2 regularization to the weights. Normalization techniques have also been utilized for regularization by researchers such as (Ba et al., 2016; Salimans and Kingma, 2016; Wu and He, 2018). (Hochreiter and Schmidhuber, 1995; Poole et al., 2014) found that adding noise can have a regularization effect. Label smoothing, a simple regularization technique particularly effective in the presence of noisy labels, has been explored by (Müller et al., 2019; Zhang et al., 2021; Li et al., 2020). Adversarial training, as proposed by (Goodfellow et al., 2015; Zhu et al., 2020; Ni et al., 2023, 2022a,b) has shown significant improvement in model performance, but it comes at the cost of increased computational effort. Dropout and its derivatives, including Adaptive Dropout by (Wan et al., 2013; Ba and Frey, 2013; Srivastava et al., 2014; Ni and Kao, 2023) have gained popularity due to their effectiveness and compatibility with other regularization techniques. Dropout enables the generation of sub-models with exponentially shared parameters during training, providing powerful regularization capabilities.

2.2. Knowledge Distillation

The concept of minimizing the output or parameter distribution between two models is commonly referred to as knowledge distillation (Hinton et al.; Furlanello et al., 2018). In knowledge distillation, a teacher model and a student model are typically employed, where the student model learns from both the ground truth labels and the teacher model during training. The teacher model serves as a guide for the student model, allowing it to learn the parameters and output distribution of the teacher model. This process can be viewed as the student model distilling knowledge from the teacher model. In the case of R-Drop (Wu et al., 2021), the generated sub-models can be seen as reciprocal teacher and student models, similar to the concept of self-distillation (Mobahi et al., 2020; Zhang et al., 2019; Zhang and Sabuncu, 2020). However, R-Drop only

applies self-distillation to the output of the model, without considering the internal representations. On the other hand, our proposed LR-Drop incorporates a *layer-wise self-distillation* approach, similar to the knowledge distillation technique employed in TinyBERT (Jiao et al., 2020). This allows for a more comprehensive knowledge interaction between the sub-models within our LR-Drop framework. It is important to note that while knowledge distillation is typically used to compress models, the primary objective of LR-Drop is to facilitate mutual learning among the sub-models within the larger model, thereby enhancing overall model performance.

3. Method: LR-Drop

This section presents a novel regularization method called LR-Drop, specifically designed for Transformer-based language models. The LR-Drop technique is illustrated in Figure 1, where it is applied to the Transformer-based model. The process begins by inputting a sample x into the model with dropout applied twice, resulting in two output distributions denoted as P_1 and P_2 . Subsequently, the *cross-entropy* loss is calculated using P_1 , P_2 , and the hard label y :

$$\mathcal{L}_{CE} = -\log_1^w(y|x) - \log_2^w(y|x). \quad (1)$$

In addition to the losses obtained from the computation with labels, LR-Drop contains three regularization losses, which are Transformer-layer regularization (containing (1) hidden states regularization loss and (2) multi-head attention regularization loss) and (3) output regularization loss. Next we will describe each of these three regularization processes and details.

3.1. Transformer-layer Regularization

As shown in Figure 1, the red box is a Transformer-layer regularization, and we regularize between the two sub-Transformer-layers sampled in each layer of the model. The two sub-models obtained by dropout random sampling are mutually teacher and student. Therefore, Transformer-layer regularization can also be seen as Transformer-layer self-distillation. The right side of Figure 1 shows a concrete representation of the Transformer-layer regularization, which contains hidden states regularization and multi-head attention regularization.

Hidden States Regularization. A fully connected feed-forward network is included in each transformer layer, which is expressed as follows:

$$\mathbf{HS}(x) = \max(0, x\mathbf{W}_1 + b_1)\mathbf{W}_2 + b_2, \quad (2)$$

where there are two linear transformations and one ReLU activation in each feed-forward network.

We regularize the knowledge from the Transformer layer outputs of the two sub-models with the following objectives:

$$\mathcal{L}_{HSR} = \text{MSE}(\mathbf{HS}(x)_1, \mathbf{HS}(x)_2), \quad (3)$$

where the matrices $\mathbf{HS}(x)_1 \in \mathbb{R}^{l \times d}$ and $\mathbf{HS}(x)_2 \in \mathbb{R}^{l \times d}$ are the hidden states of first sub-model and the second sub-model respectively, which are calculated by Equation 2. $\text{MSE}()$ is the *mean squared error* loss function, l is the input text length, and d is the hidden sizes of this two sub-models.

Multi-head Attention Regularization. The key point for the Transformer-based model to work well is that each Transformer contains a multi-headed attention module whose attention function is computed depending on the query, key and value, represented as matrices \mathbf{Q} , \mathbf{K} and \mathbf{V} . The attention function is called ‘‘Scaled Dot-Product Attention’’, it can be expressed as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{QK}^T}{\sqrt{d_k}}\right)\mathbf{V}, \quad (4)$$

where d_k is the dimension of queries and keys. Then the dot product of the query and all the keys is obtained by dividing each key by $\sqrt{d_k}$ and applying the $\text{softmax}()$ function to obtain the weights of the values. Multi-head attention is concatenated by the attention of independent initialization weights in equation (4), which enables the model to focus jointly on information from different representation subspaces. It can be expressed as:

$$\text{MHA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Concat}(A_{head}^1, \dots, A_{head}^h)\mathbf{W}^o, \quad (5)$$

where h is the number of attention heads, and A_{head}^1 denotes the i -th attention head, which is calculated by the equation (4). The matrix \mathbf{W}^o acts as a linear transformation.

Multi-head attention learns a lot of linguistic knowledge during training, and it is necessary to regularize it. Therefore, we propose multi-head attention regularization to encourage mutual learning of attention weights between two sub-models. The optimization objective is defined as:

$$\mathcal{L}_{MHAR} = \frac{1}{h} \sum_{i=1}^h \text{MSE}(\mathbf{A}_i^1, \mathbf{A}_i^2), \quad (6)$$

where h is the number of attention heads, $\mathbf{A}_i^1 \in \mathbb{R}^{l \times l}$ and $\mathbf{A}_i^2 \in \mathbb{R}^{l \times l}$ refer to the attention matrix corresponding to the i -th head of first sub-model and the second sub-model, l is the input text length, and $\text{MSE}()$ refer to the *mean squared error* loss function.

3.2. Output Regularization

In addition to regularizing the Transformer layer within the model, we also apply regularization to

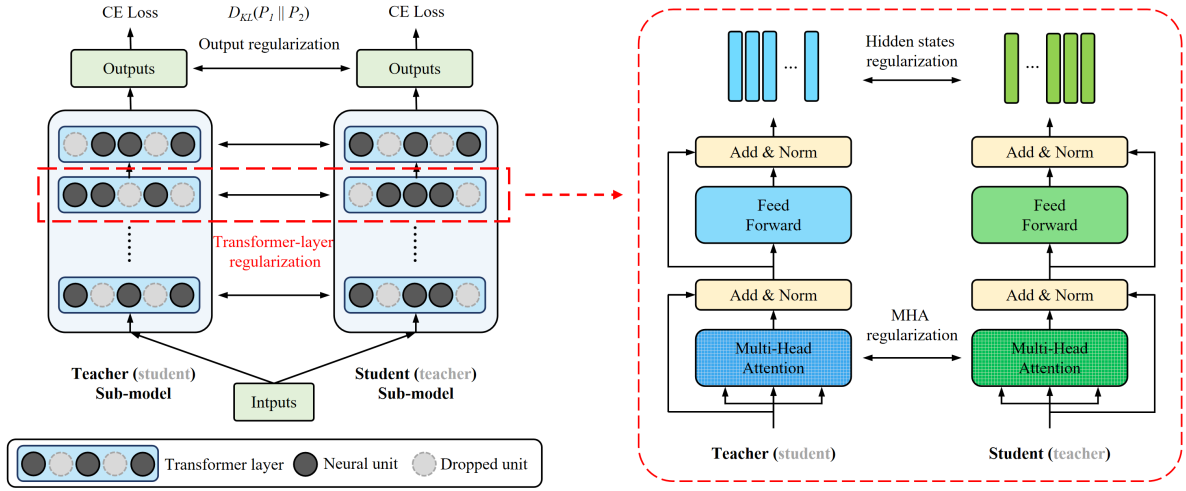


Figure 1: The proposed LR-Drop to regularize Transformer-based PLM. The left figure shows that one input will go through the two different sub-models produced by dropout twice and obtain two distributions P_1 and P_2 . The right one shows a Transformer-layer regularization containing hidden states regularization MHA regularization.

the output of the model, similar to R-Drop. Specifically, our approach, LR-Drop, aims to minimize the bidirectional KL-divergence between the output distributions of the two sub-models obtained through dropout sampling. The optimization objective is defined as:

$$\mathcal{L}_{OR} = \frac{1}{2}[\text{KL}(P_1, P_2) + \text{KL}(P_2, P_1)], \quad (7)$$

where P_1 and P_2 are the output distributions of the first and second sub-models, respectively, and $\text{KL}()$ denotes the KL-divergence loss function.

3.3. Total Optimization Objective

To summarize, the total optimization objective of our proposed LR-Drop during training is expressed as follows:

$$\mathcal{L}_{Total} = \mathcal{L}_{CE} + \alpha\mathcal{L}_{HSR} + \beta\mathcal{L}_{MHAR} + \gamma\mathcal{L}_{OR}, \quad (8)$$

where α , β , and γ are the weight coefficients for the regularization loss functions \mathcal{L}_{HSR} , \mathcal{L}_{MHAR} , and \mathcal{L}_{OR} , respectively.

4. Experiments

We assessed the effectiveness of LR-Drop on various natural language processing tasks. Our evaluation involved eight datasets for natural language understanding, six datasets for neural machine translation, and one dataset for abstractive summarization. In the table below, we use the abbreviations "RD" to refer to R-Drop and "LRD" to refer to LR-Drop in the presentation of the experimental results.

4.1. Natural Language Understanding

Datasets We begin by assessing the effectiveness of LR-Drop on natural language understanding tasks. The GLUE Benchmark consists of eight English natural language understanding tasks, which vary in domains, data volumes, and difficulty levels.

(1) **RTE** (Dagan et al., 2006; Bar Haim et al., 2006; Giampiccolo et al., 2007): This dataset comprises a series of natural language inference datasets used in annual text challenges.

(2) **MNLI** (Williams et al., 2018): In this task, a premise and a hypothesis are given, and the objective is to predict whether the premise supports or contradicts the hypothesis, or neither.

(3) **MRPC** (Dolan and Brockett, 2005): Given a pair of sentences, the task is to determine whether their semantics are the same.

(4) **STS-B** (Agirre et al., 2007): Each data instance consists of a pair of sentences along with a similarity score ranging from 1 to 5. The task involves discrete regression to predict the scores.

(5) **QQP** (Iyer et al., 2017): This task involves identifying whether a pair of questions are semantically identical.

(6) **SST-2** (Socher et al., 2013): This binary task requires predicting whether a sentence is positive or negative.

(7) **QNLI** (Rajpurkar et al., 2016): The objective of this task is to determine if a given question can be answered using the context sentence.

(8) **CoLA** (Warstadt et al., 2018): This task focuses on assessing the grammatical accuracy of a sentence.

Experimental Settings In this subsection, we

Model	MNLI	MRPC	QNLI	QQP	RTE	SST-2	STS-B	CoLA	Avg
BERT-base	83.8	85.3	90.8	91.0	68.2	92.4	89.3	62.3	82.85
BERT-base + RD	85.5	87.3	92.0	91.4	71.1	93.0	89.6	62.6	84.06
BERT-base + LRD	86.1	88.0	92.3	91.4	71.8	93.2	90.2	63.4	84.55
RoBERTa-large	90.2	90.9	94.7	92.2	86.6	96.4	92.4	68.0	88.93
RoBERTa-large + RD	90.9	91.4	95.2	92.5	88.4	96.9	92.5	70.0	89.73
RoBERTa-large + LRD	91.3	91.8	95.9	93.2	89.2	97.4	92.7	71.8	90.41
ELECTRA-large	90.9	90.8	95.0	92.4	88.0	96.9	92.6	69.1	89.46
ELECTRA-large + RD	91.2	91.3	95.6	92.6	88.9	97.4	92.8	70.5	90.03
ELECTRA-large + LRD	91.7	92.1	96.2	93.3	89.5	97.6	93.1	71.9	90.68

Table 1: Performances on natural language understanding tasks of GLUE benchmark. Significance test: the average performance of LR-Drop and R-Drop on the GLUE datasets was t-tested to obtain a p-value of $0.0034 < 0.01$.

employ three publicly available pre-trained language models (PLMs) as the baseline models for our experiments to evaluate the effectiveness of LR-Drop. The chosen PLMs are BERT-base, RoBERTa-large, and ELECTRA-large. Different tasks may require different parameter settings, so we dynamically adjust the coefficients α , β , and γ from the set $\{0.01, 0.05, 0.1, 0.5\}$ accordingly. The experimental methodology for the comparative models follows the approach outlined in previous research. For the STS-B task, we use the Pearson correlation as the evaluation metric, while for CoLA, we employ Matthew’s correlation. The remaining tasks are evaluated based on Accuracy. We report the mean results of 5 runs to ensure statistical reliability. The experiments were conducted using an RTX 3090 GPU.

Experimental Results The experimental results are presented in Table 1. When applying LR-Drop to the BERT-base model, we observed improvements in fine-tuning scores across multiple tasks. For the MNLI task, the fine-tuning score increased from 83.8 to 86.1. In the MRPC task, the score improved from 85.3 to 88.8. The QNLI task saw an improvement from 90.8 to 92.3, and the QQP task improved from 91.0 to 91.4. For the RTE task, the score increased from 68.2 to 71.8. The SST-2 task showed an improvement from 92.4 to 93.2, and the STS-B task improved from 89.3 to 90.2. Similarly, the RoBERTa-large and ELECTRA-large models also exhibited performance improvements of more than 1 point per dataset when using LR-Drop.

Across the eight datasets, the BERT-base + LR-Drop, RoBERTa-large + LR-Drop, and ELECTRA-large + LR-Drop models achieved impressive average scores of 84.55, 90.41, and 90.68, respectively. LR-Drop significantly improved the performance of the three baseline models: BERT-base, RoBERTa-large, and ELECTRA-large, by 1.70 points, 1.48 points, and 1.22 points, respectively. Furthermore, when compared to the previous method R-Drop, our proposed LR-Drop demonstrated average im-

provements of 0.49 points, 0.65 points, and 0.68 points, respectively. These results indicate that the performance of our LR-Drop regularization is particularly enhanced when applied to stronger baseline models. Additionally, the effectiveness of LR-Drop is evident across different neural language models, resulting in improved performance in natural language understanding tasks.

4.2. Neural Machine Translation

Datasets The datasets used for neural machine translation were obtained from the International Workshop on Spoken Language Translation (IWSLT) competitions. These datasets consist of translations between English and German (En \leftrightarrow De), English and Spanish (En \leftrightarrow Es), English and French (En \leftrightarrow Fr), and English and Chinese (En \leftrightarrow Zh). Specifically, we used the IWSLT14 dataset for English to German and vice versa, the IWSLT14 dataset for English to Spanish and vice versa, the IWSLT17 dataset for English to French and vice versa, and the IWSLT dataset for English to Chinese and vice versa. The IWSLT dataset contains approximately 170,000 pairs of sentences for training, 7,000 pairs for validation, and 7,000 pairs for testing. These datasets serve as valuable resources for training and evaluating our neural machine translation models.

Experimental Settings The experimental configuration outlined in (Wu et al., 2021) is followed in this study. Our benchmark model is the Transformer network proposed by (Vaswani et al., 2017). The specific configurations for the IWSLT translations are specified under the transformer_iwslt_de_en setting. To explore different settings, the coefficients α , β , and γ are dynamically varied within the set $\{0.1, 0.5, 1\}$. The implementation of our models is carried out using the Fairseq framework. For evaluating the performance of the models on neural machine translation tasks, we utilize BLEU scores. The reported results are

Model	En to De	De to En	En to Fr	Fr to En	En to Zh	Zh to En	En to Es	Es to En	Avg
Transformer	28.57	34.64	35.9	36.1	26.3	18.4	39.0	40.6	32.44
Transformer + RD	30.72	37.25	38.0	38.9	28.1	19.5	41.8	43.2	34.68
Transformer + LRD	30.95	37.87	38.8	39.6	28.6	20.3	42.6	44.1	35.35

Table 2: BLEU scores on 8 IWSLT machine translation tasks. Significance test: the average performance of LR-Drop and R-Drop on the 8 datasets was t-tested to obtain a p-value of $0.0037 < 0.01$.

the averages obtained from five trial runs to ensure robustness. The experiments are conducted on an RTX 3090 GPU, which serves as the hardware for the experiments.

Experimental Results The experimental results are presented in Table 2. When applying the Transformer model to our LR-Drop technique, we observed improvements in the fine-tuning scores for various translation tasks. For the English to German translation task, the fine-tuning score increased from 28.57 points to 30.95 points. Similarly, for the German to English translation, the score improved from 34.64 points to 37.87 points. For the English to French translation, the score improved from 35.9 points to 38.8 points, and for the French to English translation, the score improved from 36.1 points to 39.6 points. In the English to Chinese translation task, the score improved from 26.3 points to 28.6 points, and in the Chinese to English translation, the score improved from 18.4 points to 20.3 points. Additionally, for the English to Spanish translation, the score improved from 39.0 points to 42.6 points, and for the Spanish to English translation, the score improved from 40.6 points to 44.1 points.

Comparing the Transformer model with the LR-Drop technique to the original Transformer, we observed an average improvement of 2.91 points across the eight IWSLT machine translation tasks. Furthermore, the Transformer model with R-Drop achieved an average score of 34.68, while our Transformer model with LR-Drop achieved an average score of 35.35. These results demonstrate the effectiveness of our method in improving the performance of neural machine translation tasks.

4.3. Abstractive Summarization

Datasets The abstract summarization task aims to condense lengthy sentences or documents into concise sequences while preserving the main content. It is a generation task. To evaluate the effectiveness of LR-Drop on the abstract summarization task, we employ the CNN/Daily Mail dataset. This dataset consists of news documents (source) and their corresponding highlighted summaries (target), which are extracted from CNN and Daily Mail websites. The dataset comprises 287,226 training documents, 13,368 validation documents, and 11,490 test documents. We preprocess the dataset follow-

Model	RG-1	RG-2	RG-L
Transformer ^a	39.50	16.06	36.63
ProphetNet ^b	44.02	21.17	41.30
BART ^c	44.16	21.28	40.90
PEGASUS ^d	44.17	21.47	41.11
BART + R3F ^e	44.38	21.53	41.17
BART + RD ^f	44.51	21.58	41.24
BART + LRD	44.58	21.63	41.30

Table 3: The ROUGE scores, consisting of ROUGE-1, ROUGE-2, and ROUGE-L, are presented for the CNN/Daily Mail summarization dataset. A significance test was carried out, comparing the mean performance of LR-Drop and R-Drop. A t-test yielded a p-value of 0.0064, which is less than 0.01. References: ^a: (Vaswani et al., 2017); ^b: (Qi et al., 2020); ^c: (Lewis et al., 2020); ^d: (Zhang et al., 2020); ^e: (Aghajanyan et al., 2020); ^f: (Wu et al., 2021).

ing the guidelines provided by Wu et al. (2021) (Wu et al., 2021).

Experimental Settings In this subsection, we conduct experiments using the BART pre-training model as a baseline to assess the efficacy of our proposed LR-Drop regularization technique. During the fine-tuning process, we apply the BART model with LR-Drop. The coefficients α , β , and γ are set to 0.1, 0.2, and 0.5, respectively. The remaining hyper-parameter settings are consistent with the original paper on the BART model. To evaluate the performance of the models on the abstract summarization task, we employ the ROUGE F1 score. The reported results are the mean scores obtained from five independent runs. The experiments are conducted on an RTX 3090 GPU.

Experimental Results To provide a comprehensive evaluation of the different methods, we report the ROUGE-1 (RG-1) and two-sequence ROUGE-2 (RG-2) overlaps to assess the information content, as well as the longest common subsequence ROUGE-L (RG-L) scores to evaluate fluency. The experimental results are presented in Table 3. Notably, BART, when trained with our LR-Drop technique, outperforms all other models and achieves state-of-the-art performance. Compared to the original BART model, BART + LR-Drop demonstrates improvements of 0.41 points in RG-1, 0.35 points in RG-2, and 0.40 points in RG-L scores. Addition-

Model	MNLI	MRPC	QNLI	QQP	RTE	SST-2	STS-B	CoLA	Avg
LR-Drop	86.1	88.0	92.3	91.4	71.8	93.2	90.2	63.4	84.55
– w/o \mathcal{L}_{HSR}	85.7	87.4	92.2	91.4	71.6	93.1	89.6	62.7	84.21
– w/o \mathcal{L}_{MHAR}	85.9	87.4	92.2	91.6	71.7	93.1	89.9	62.7	84.30
– w/o \mathcal{L}_{OR}	85.8	87.5	92.1	91.5	71.5	93.2	89.7	62.8	84.26

Table 4: Ablation study of LR-Drop.

Methods	size of the training set					
	1K	2K	4K	8K	16K	32K
BERT-base	53.10	54.24	57.34	62.39	86.58	89.16
BERT-base + R-Drop	54.12	56.21	57.87	62.74	87.10	89.56
BERT-base + LR-Drop	55.04	56.53	58.56	63.47	88.03	90.07

Table 5: LR-Drop performance on training sets of different sizes. Dataset: SST-2, baseline model: BERT-base.

ally, BERT + LR-Drop surpasses BERT + R3F and BERT + R-Drop. These experimental findings highlight the effectiveness of our LR-Drop technique in enhancing the performance of abstract summarization tasks.

5. Study and Analysis on LR-Drop

In this section, we present a comprehensive investigation and analysis of LR-Drop, considering multiple aspects such as ablation study, training set sizes, dropout times, and loss landscape analysis. All experiments in this section are conducted using the datasets from the General Language Understanding Evaluation (GLUE) benchmark.

5.1. Ablation Study

To further analyze the role of each component in LR-Drop, we conducted ablation experiments by removing individual parts of the loss function: \mathcal{L}_{HSR} , \mathcal{L}_{MHAR} , and \mathcal{L}_{OR} . This allowed us to assess the contribution of each component to the overall performance of LR-Drop. The following three ablated methods were obtained: (1) LR-Drop w/o \mathcal{L}_{HSR} : This method removes the hidden states regularization component from LR-Drop. (2) LR-Drop w/o \mathcal{L}_{MHAR} : This method removes the multi-head attention regularization component from LR-Drop. (3) LR-Drop w/o \mathcal{L}_{OR} : This method removes the output regularization component from LR-Drop.

Table 4 presents the results of the ablation experiments. It is evident that the full LR-Drop method achieves the best performance. When the hidden states regularization component is ablated, the average score of LR-Drop on GLUE decreases to 84.21, which is 0.34 points lower than the full LR-Drop. Similarly, when the multi-head attention regularization component is removed, the average score of LR-Drop on GLUE decreases to 84.30, which is 0.25 points lower than the full LR-Drop. Lastly,

when the output regularization component is ablated, the average score of LR-Drop on GLUE decreases to 84.26, which is 0.29 points lower than the full LR-Drop.

These results highlight the importance of each component in LR-Drop and demonstrate that the combination of all components leads to the highest performance on GLUE tasks.

5.2. LR-Drop Performance on Training Sets of Different Sizes

This study also investigates the performance of LR-Drop across training datasets of varying sizes. To examine this, we partitioned the SST-2 training set into different subsets and utilized the BERT-base model with the same structure as previously described. The results of this experiment are presented in Table 5, where we compare LR-Drop with R-Drop. As shown in Table 5, our approach, LR-Drop, not only achieves superior performance on large training datasets but also demonstrates significant improvements on smaller datasets. In comparison to the BERT-base model, our enhanced model, BERT-base + LR-Drop, achieves performance boosts of 1% to 2% across distinct training sets. Furthermore, LR-Drop outperforms R-Drop on training sets of various sizes. These experimental findings highlight the effectiveness of LR-Drop in providing robust regularization for the model, even when data availability is limited.

5.3. 3-time LR-Drop

From the above, it is evident that LR-Drop is based on the concept of "twice dropout" to generate two sub-networks for regularization during training. In our research, we extended this approach by increasing the number of forward propagations from two to three, meaning that each example passes through the model three times in LR-Drop. During

Model	MNLI	MRPC	QNLI	QQP	RTE	SST-2	STS-B	CoLA	Avg
LR-Drop-2	86.1	88.0	92.3	91.4	71.8	93.2	90.2	63.4	84.55
LR-Drop-3	86.2	88.1	92.3	91.3	71.8	93.4	90.1	63.4	84.58

Table 6: The effect of the number of LR-Drop samples on performance.




Figure 2: 2D (left) and 3D (right) visualization of loss function minima selected by BERT-base with standard training (ST) and LR-Drop on SST-2 dataset.

training, we also incorporate hidden states regularization, multi-head attention regularization, and output regularization to facilitate mutual learning and knowledge exchange among the three sub-models. The experimental results, as presented in Table 6, demonstrate that LR-Drop-3, which consists of three sub-models, yields a slight improvement of 0.03 points over LR-Drop-2, the original LR-Drop with two sub-models. However, it is worth noting that increasing the number of dropout samples tends to increase computational consumption. Therefore, we believe that utilizing LR-Drop-2 regularization twice is sufficient for achieving effective model training. However, if computational resources are not a constraint, one can consider using LR-Drop-3.

It is important to mention that an increase in iterations corresponds to an increase in resource usage. Therefore, we are of the view that the double iteration approach in LR-Drop-2 provides adequate regularization during model training. However, if computational resources are not a limitation, one might explore the option of deploying LR-Drop-3.

5.4. Loss Landscape Analysis

To provide a more visual analysis of the regularization effect of LR-Drop, we employed a method proposed by Li et al. (2018) to visualize high-dimensional non-convex loss functions. This method allowed us to graphically depict the loss landscapes surrounding the empirical risk minima obtained through standard training and LR-Drop training, both with the same model structure. Figure 2 presents 2D and 3D visualizations of these loss landscapes.

In this visualization, we considered the dimen-

sions of θ and randomly extracted two direction-representing vectors, d_x and d_y , from a Gaussian distribution with a mean of zero. The scale of these vectors was set to be equivalent to the variance of the weight values. We then introduced linear perturbations α and β to a central point θ^* , resulting in a loss function that represents variations in these two random directions:

$$f(d_x, d_y) = \mathcal{L}(\theta^* + \alpha d_x + \beta d_y) \quad (9)$$

The analysis points out that the proposed LR-Drop indeed selects flatter loss landscapes by dynamically creating perturbation. Numerous investigations have evidenced that a flatter loss landscape generally implies superior generalization (Keskar et al., 2019; Ishida et al., 2020).

6. Conclusion

In this work, we have introduced LR-Drop, a simple yet effective regularization technique based on dropout. LR-Drop utilizes layer-wise self-distillation between two sub-models generated by dropout to improve model performance. During training, LR-Drop regularizes different representations in Transformer-based Pre-trained Language Models (PLMs), including the hidden state, the multi-headed attention matrix, and the output distribution of the prediction layer. Importantly, LR-Drop does not introduce additional model parameters or modify the original architecture of the language model, making it applicable to various Transformer-based PLMs. Through extensive experiments on a diverse set of 15 datasets, including natural language understanding, neural machine translation, and abstractive summarization tasks, we have

demonstrated that LR-Drop achieves superior performance compared to existing methods. Our results even surpass state-of-the-art performance in these tasks. This highlights the effectiveness and versatility of LR-Drop as a regularization technique for Transformer-based PLMs.

7. Acknowledgements

This work is supported by Open Research Fund of Key Laboratory of Intelligent Education Technology and Application of Zhejiang Province, Zhejiang Normal University (jykf21002w), National Key Research and Development Program of China (2022YFF0902100), Postdoctoral Fellowship Program of CPSF (GZC20232873), Guangdong Basic and Applied Basic Research Foundation (2023A1515110718), National Natural Science Foundation of China (62376262), Shenzhen Science and Technology Innovation Program (KQTD20190929172835662), Shenzhen Basic Research Foundation (JCYJ20210324115614039 and JCYJ20200109113441941).

8. Bibliographical References

- Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2020. Better fine-tuning by reducing representational collapse. In *International Conference on Learning Representations*.
- Eneko Agirre, Lluís M’arquez, and Richard Wicentowski, editors. 2007. *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*. Association for Computational Linguistics.
- Jimmy Ba and Brendan Frey. 2013. Adaptive dropout for training deep neural networks. *Advances in neural information processing systems*, 26:3084–3092.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Roy Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. 2006. The second PASCAL recognising textual entailment challenge.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019a. What does bert look at? an analysis of bert’s attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2019b. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. The PASCAL recognising textual entailment challenge. In *Machine learning challenges. evaluating predictive uncertainty, visual object classification, and recognising textual entailment*. Springer.
- William B Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the International Workshop on Paraphrasing*.
- Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. 2018. Born again neural networks. In *International Conference on Machine Learning*, pages 1607–1616. 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*.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. *ICLR*.
- Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network.
- Sepp Hochreiter and Jürgen Schmidhuber. 1995. Simplifying neural nets by discovering flat minima. In *Advances in neural information processing systems*, pages 529–536.
- Takashi Ishida, Ikko Yamane, Tomoya Sakai, Gang Niu, and Masashi Sugiyama. 2020. Do we need zero training loss after achieving zero training error? In *International Conference on Machine Learning*, pages 4604–4614. PMLR.
- Shankar Iyer, Nikhil Dandekar, and Kornl Csernai. 2017. [First quora dataset release: Question pairs](#).
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. Tinybert: Distilling bert for natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4163–4174.
- Guoliang Kang, Jun Li, and Dacheng Tao. 2016. Shakeout: A new regularized deep neural network training scheme. In *Thirtieth AAAI Conference on Artificial Intelligence*.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Nitish Shirish Keskar, Jorge Nocedal, Ping Tak Peter Tang, Dheevatsa Mudigere, and Mikhail Smelyanskiy. 2019. On large-batch training for deep learning: Generalization gap and sharp minima. In *5th International Conference on Learning Representations, ICLR 2017*.
- Anders Krogh and John A Hertz. 1992. A simple weight decay can improve generalization. In *Advances in neural information processing systems*, pages 950–957.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Weizhi Li, Gautam Dasarathy, and Visar Berisha. 2020. Regularization via structural label smoothing. In *International Conference on Artificial Intelligence and Statistics*, pages 1453–1463. PMLR.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Xuezhe Ma, Yingkai Gao, Zhiting Hu, Yaoliang Yu, Yuntian Deng, and Eduard Hovy. 2016. Dropout with expectation-linear regularization. *arXiv preprint arXiv:1609.08017*.
- Hossein Mobahi, Mehrdad Farajtabar, and Peter Bartlett. 2020. Self-distillation amplifies regularization in hilbert space. *Advances in Neural Information Processing Systems*, 33:3351–3361.
- Rafael Müller, Simon Kornblith, and Geoffrey Hinton. 2019. When does label smoothing help? *arXiv preprint arXiv:1906.02629*.
- Shiwen Ni and Hung-Yu Kao. 2023. Masked siamese prompt tuning for few-shot natural language understanding. *IEEE Transactions on Artificial Intelligence*.
- Shiwen Ni, Jiawen Li, and Hung-Yu Kao. 2022a. Hat4rd: Hierarchical adversarial training for rumor detection in social media. *Sensors*, 22(17):6652.
- Shiwen Ni, Jiawen Li, and Hung-Yu Kao. 2022b. R-at: Regularized adversarial training for natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6427–6440.
- Shiwen Ni, Jiawen Li, Min Yang, and Hung-Yu Kao. 2023. Dropattack: A random dropped weight attack adversarial training for natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Ben Poole, Jascha Sohl-Dickstein, and Surya Ganguli. 2014. Analyzing noise in autoencoders and deep networks. *arXiv preprint arXiv:1406.1831*.
- Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. Prophetnet: Predicting future n-gram for sequence-to-sequence pre-training. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2401–2410.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *EMNLP*.
- Tim Salimans and Durk P Kingma. 2016. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems*, 29:901–909.
- 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 *EMNLP*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. 2013. Regularization of neural networks using dropconnect. In *International conference on machine learning*, pages 1058–1066. PMLR.

- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2018. Neural network acceptability judgments. *arXiv preprint 1805.12471*.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL*.
- Lijun Wu, Juntao Li, Yue Wang, Qi Meng, Tao Qin, Wei Chen, Min Zhang, Tie-Yan Liu, et al. 2021. R-drop: Regularized dropout for neural networks. *Advances in Neural Information Processing Systems*, 34:10890–10905.
- Yuxin Wu and Kaiming He. 2018. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pre-training for language understanding. *Advances in neural information processing systems*, 32.
- Chang-Bin Zhang, Peng-Tao Jiang, Qibin Hou, Yunchao Wei, Qi Han, Zhen Li, and Ming-Ming Cheng. 2021. Delving deep into label smoothing. *IEEE Transactions on Image Processing*, 30:5984–5996.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. 2019. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3713–3722.
- Zhilu Zhang and Mert Sabuncu. 2020. Self-distillation as instance-specific label smoothing. *Advances in Neural Information Processing Systems*, 33:2184–2195.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2020. Freelib: Enhanced adversarial training for natural language understanding. In *ICLR*.
- Konrad Zolna, Devansh Arpit, Dendi Suhubdy, and Yoshua Bengio. 2018. Fraternal dropout. In *International Conference on Learning Representations*.