

# Self-supervised Regularization for Text Classification

Meng Zhou\*

Shanghai Jiao Tong University,  
China

zhoumeng9904@sjtu.edu.cn

Zechen Li\*

Northeastern University,  
United States

li.zec@northeastern.edu

Pengtao Xie<sup>†</sup>

UC San Diego, United States  
plxie@eng.ucsd.edu

## Abstract

Text classification is a widely studied problem and has broad applications. In many real-world problems, the number of texts for training classification models is limited, which renders these models prone to overfitting. To address this problem, we propose SSL-Reg, a data-dependent regularization approach based on self-supervised learning (SSL). SSL (Devlin et al., 2019a) is an unsupervised learning approach that defines auxiliary tasks on input data without using any human-provided labels and learns data representations by solving these auxiliary tasks. In SSL-Reg, a supervised classification task and an unsupervised SSL task are performed simultaneously. The SSL task is unsupervised, which is defined purely on input texts without using any human-provided labels. Training a model using an SSL task can prevent the model from being overfitted to a limited number of class labels in the classification task. Experiments on 17 text classification datasets demonstrate the effectiveness of our proposed method. Code is available at <https://github.com/UCSD-AI4H/SSReg>.

## 1 Introduction

Text classification (Korde and Mahender, 2012; Lai et al., 2015; Wang et al., 2017; Howard and Ruder, 2018) is a widely studied problem in natural language processing and finds broad applications. For example, given clinical notes of a patient, judge whether this patient has heart diseases. Given a scientific paper, judge whether it is about NLP. In many real-world text classification problems, texts available for training are oftentimes limited. For instance, it is difficult to obtain many clinical notes from hospitals due

to concern of patient privacy. It is well known that when training data is limited, models tend to overfit to training data and perform less well on test data.

To address overfitting problems in text classification, we propose a data-dependent regularizer called SSL-Reg based on self-supervised learning (SSL) (Devlin et al., 2019a; He et al., 2019; Chen et al., 2020) and use it to regularize the training of text classification models, where a supervised classification task and an unsupervised SSL task are performed simultaneously. SSL (Devlin et al., 2019a; He et al., 2019; Chen et al., 2020) is an unsupervised learning approach that defines auxiliary tasks on input data without using any human-provided labels and learns data representations by solving these auxiliary tasks. For example, BERT (Devlin et al., 2019a) is a typical SSL approach where an auxiliary task is defined to predict masked tokens and a text encoder is learned by solving this task. In existing SSL approaches for NLP, an SSL task and a target task are performed sequentially. A text encoder is first trained by solving an SSL task defined on a large collection of unlabeled texts. Then this encoder is used to initialize an encoder in a target task. The encoder is finetuned by solving the target task. A potential drawback of performing SSL task and target task sequentially is that text encoder learned in SSL task may be overridden after being finetuned in target task. If training data in the target task is small, the finetuned encoder has a high risk of being overfitted to training data.

To address this problem, in SSL-Reg we perform SSL task and target tasks (which is classification) simultaneously. In SSL-Reg, an SSL loss serves as a regularization term and is optimized jointly with a classification loss. SSL-Reg enforces a text encoder to jointly solve two tasks: an unsupervised SSL task and a supervised text classification task. Due to the presence of the SSL task, models are less likely to be biased to the classification task defined on small-sized training

\*Equal contribution.

<sup>†</sup>Corresponding author.

data. We perform experiments on 17 datasets, where experimental results demonstrate the effectiveness of SSL-Reg in alleviating overfitting and improving generalization performance.

The major contributions of this paper are:

- We propose SSL-Reg, which is a data-dependent regularizer based on SSL, to reduce the risk that a text encoder is biased to a data-deficient classification task on small-sized training data.
- Experiments on 17 datasets demonstrate the effectiveness of our approaches.

The rest of this paper is organized as follows. Section 2 reviews related works. Section 3 and 4 present methods and experiments. Section 5 concludes the paper and discusses future work.

## 2 Related Works

### 2.1 Self-supervised Learning for NLP

SSL aims to learn meaningful representations of input data without using human annotations. It creates auxiliary tasks solely using input data and forces deep networks to learn highly effective latent features by solving these auxiliary tasks. In NLP, various auxiliary tasks have been proposed for SSL, such as next token prediction in GPT (Radford et al.), masked token prediction in BERT (Devlin et al., 2019a), text denoising in BART (Lewis et al., 2019), contrastive learning (Fang et al., 2020), and so on. These models have achieved substantial success in learning language representations. The GPT model (Radford et al.) is a language model based on Transformer (Vaswani et al., 2017). Different from Transformer, which defines a conditional probability on an output sequence given an input sequence, GPT defines a marginal probability on a single sequence. In GPT, conditional probability of the next token given a historical sequence is defined using a Transformer decoder. Weight parameters are learned by maximizing likelihood on token sequences. BERT (Devlin et al., 2019a) aims to learn a Transformer encoder for representing texts. BERT’s model architecture is a multi-layer bidirectional Transformer encoder. In BERT, Transformer uses bidirectional self-attention. To train the encoder, BERT masks some percentage of input tokens at random, and then predicts those masked tokens by feeding hidden vectors (produced by the encoder)

corresponding to masked tokens into an output softmax over word vocabulary. BERT-GPT (Wu et al., 2019) is a model used for sequence-to-sequence modeling where a pretrained BERT is used to encode input text and GPT is used to generate output texts. In BERT-GPT, pretraining of BERT encoder and GPT decoder is conducted separately, which may lead to inferior performance. Auto-Regressive Transformers (BART) (Lewis et al., 2019) has a similar architecture as BERT-GPT, but trains BERT encoder and GPT decoder jointly. To pretrain BART weights, input texts are corrupted randomly, such as token masking, token deletion, text infilling, and so forth, then a network is learned to reconstruct original texts. ALBERT (Lan et al., 2019) uses parameter-reduction methods to reduce memory consumption and increase training speed of BERT. It also introduces a self-supervised loss that models inter-sentence coherence. RoBERTa (Liu et al., 2019a) is a replication study of BERT pretraining. It shows that BERT’s performance can be greatly improved by carefully tuning training processes, such as (1) training models longer, with larger batches, over more data; (2) removing the next sentence prediction objective; (3) training on longer sequences, and so on. XLNet (Yang et al., 2019) learns bidirectional contexts by maximizing expected likelihood over all permutations of factorization order and uses a generalized autoregressive pretraining mechanism to overcome the pretrain-finetune discrepancy of BERT. T5 (Raffel et al., 2019) compared pretraining objectives, architectures, unlabeled datasets, transfer approaches on a wide range of language understanding tasks and proposed a unified framework that casts these tasks as a text-to-text task. ERNIE 2.0 (Sun et al., 2019) proposed a continual pretraining framework which builds and learns incrementally pretraining tasks through constant multi-task learning, to capture lexical, syntactic and semantic information from training corpora. Gururangan et al. (2020) proposed task adaptive pretraining (TAPT) and domain adaptive pretraining (DAPT). Given a RoBERTa model pretrained on large-scale corpora, TAPT continues to pretrain RoBERTa on training dataset of target task. DAPT continues to pretrain RoBERTa on datasets that have small domain differences with data in target tasks. The difference between our proposed SSL-Reg method with TAPT and DAPT is that SSL-Reg uses a self-supervised task (e.g., mask token prediction) to regularize the finetuning of

RoBERTa where text classification task and self-supervised task are performed jointly. In contrast, TAPT and DAPT use self-supervised task for pretraining, where text classification task and self-supervised task are performed sequentially. The connection between our method and TAPT is that they both leverage texts in target tasks to perform self-supervised learning, in addition to SSL on large-scale external corpora. Different from SSL-Reg and TAPT, DAPT uses domain-similar texts rather than target texts for additional SSL.

## 2.2 Self-supervised Learning in General

Self-supervised learning has been widely applied to other application domains, such as image classification (He et al., 2019; Chen et al., 2020), graph classification (Zeng and Xie, 2021), visual question answering (He et al., 2020a), and so forth, where various strategies have been proposed to construct auxiliary tasks, based on temporal correspondence (Li et al., 2019; Wang et al., 2019a), cross-modal consistency (Wang et al., 2019b), rotation prediction (Gidaris et al., 2018; Sun et al., 2020), image inpainting (Pathak et al., 2016), automatic colorization (Zhang et al., 2016), context prediction (Nathan Mundhenk et al., 2018), and so on. Some recent works studied self-supervised representation learning based on instance discrimination (Wu et al., 2018) with contrastive learning. Oord et al. (2018) proposed contrastive predictive coding, which predicts the future in latent space by using powerful autoregressive models, to extract useful representations from high-dimensional data. Bachman et al. (2019) proposed a self-supervised representation learning approach based on maximizing mutual information between features extracted from multiple views of a shared context. MoCo (He et al., 2019) and SimCLR (Chen et al., 2020) learned image encoders by predicting whether two augmented images were created from the same original image. Srinivas et al. (2020) proposed to learn contrastive unsupervised representations for reinforcement learning. Khosla et al. (2020) investigated supervised contrastive learning, where clusters of points belonging to the same class were pulled together in embedding space, while clusters of samples from different classes were pushed apart. Klein and Nabi (2020) proposed a contrastive self-supervised learning approach for commonsense reasoning. He et al. (2020b);

Yang et al. (2020) proposed an Self-Trans approach which applied contrastive self-supervised learning on top of networks pretrained by transfer learning.

Compared with supervised learning that requires each data example to be labeled by humans or semi-supervised learning which requires part of data examples to be labeled, self-supervised learning is similar to unsupervised learning because it does not need human-provided labels. The key difference between SSL and unsupervised learning is that SSL focuses on learning data representations by solving auxiliary tasks defined on un-labeled data while unsupervised learning is more general and aims to discover latent structures from data, such as clustering, dimension reduction, manifold embedding (Roweis and Saul, 2000), and so on.

## 2.3 Text Classification

Text classification (Minaee et al., 2020) is one of the key tasks in natural language processing and has a wide range of applications, such as sentiment analysis, spam detection, tag suggestion, and so forth. A number of approaches have been proposed for text classification. Many of them are based on RNNs. Liu et al. (2016) use multi-task learning to train RNNs, utilizing the correlation between tasks to improve text classification performance. Tai et al. (2015) generalize sequential LSTM to tree-structured LSTM to capture the syntax of sentences for achieving better classification performance. Compared with RNN-based models, CNN-based models are good at capturing local and position-invariant patterns. Kalchbrenner et al. (2014) proposed dynamic CNN, which uses dynamic k-max-pooling to explicitly capture short-term and long-range relations of words and phrases. Zhang et al. (2015) proposed a character-level CNN model for text classification, which can deal with out-of-vocabulary words. Hybrid methods combine RNN and CNN to explore the advantages of both. Zhou et al. (2015) proposed a convolutional LSTM network, which uses a CNN to extract phrase-level representations, then feeds them to an LSTM network to represent the whole sentence.

## 3 Methods

To alleviate overfitting in text classification, we propose SSL-Reg, which is a regularization

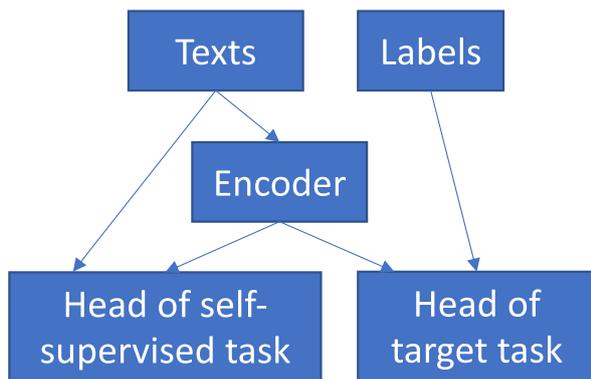


Figure 1: Illustration of SSL-Reg. Input texts are fed into a text encoder. Encodings of these texts and their corresponding labels are fed into the head of a target task (e.g., classification) which yields a classification loss. In a self-supervised task, inputs are encodings of texts and outputs are constructed on original texts (e.g., masked tokens). The classification task and SSL task share the same text encoder and losses of these two tasks are optimized jointly to learn the text encoder.

approach based on self-supervised learning (SSL), where an unsupervised SSL task and a supervised text classification task are performed jointly.

### 3.1 SSL-based Regularization

SSL-Reg uses a self-supervised learning task to regularize a text classification model. Figure 1 presents an illustration of SSL-Reg. Given training texts, we encode them using a text encoder. Then on top of text encodings, two tasks are defined. One is a classification task, which takes the encoding of a text as input and predicts the class label of this text. Prediction is conducted using a classification head. The other task is SSL. The loss of the SSL task serves as a data-dependent regularizer to alleviate overfitting. The SSL task has a predictive head. These two tasks share the same text encoder. Formally, SSL-Reg solves the following optimization problem:

$$\mathcal{L}^{(c)}(D, L; \mathbf{W}^{(e)}, \mathbf{W}^{(c)}) + \lambda \mathcal{L}^{(p)}(D, \mathbf{W}^{(e)}, \mathbf{W}^{(p)}) \quad (1)$$

where  $D$  represents training texts and  $L$  represents their labels.  $\mathbf{W}^{(e)}$ ,  $\mathbf{W}^{(c)}$ , and  $\mathbf{W}^{(p)}$  denote text encoder, classification head in classification task, and prediction head in SSL task, respectively.  $\mathcal{L}^{(c)}$  denotes classification loss and  $\mathcal{L}^{(p)}$  denotes SSL loss.  $\lambda$  is a tradeoff parameter.

At the core of SSL-Reg is using SSL to learn a text encoder that is robust to overfitting. Our methods can be used to learn any text encoder. In this

work, we perform the study using a Transformer encoder, while noting that other text encoders are also applicable.

### 3.2 Self-supervised Learning Tasks

In this work, we use two self-supervised learning tasks—masked token prediction (MTP) and sentence augmentation type prediction (SATP)—to perform our studies while noting that other SSL tasks are also applicable.

- **Masked Token Prediction (MTP)** This task is used in BERT. Some percentage of input tokens are masked at random. Texts with masked tokens are fed into a text encoder that learns a latent representation for each token including the masked ones. The task is to predict these masked tokens by feeding hidden vectors (produced by the encoder) corresponding to masked tokens into an output softmax over word vocabulary.

- **Sentence Augmentation Type Prediction (SATP)** Given an original text  $o$ , we apply different types of augmentation methods to create augmented texts from  $o$ . We train a model to predict which type of augmentation was applied to an augmented text. We consider four types of augmentation operations used in Wei and Zou (2019), including synonym replacement, random insertion, random swap, and random deletion. Synonym replacement randomly chooses 10% of non-stop tokens from original texts and replaces each of them with a randomly selected synonym. In random insertion, for a randomly chosen non-stop token in a text, among the synonyms of this token, one randomly selected synonym is inserted into a random position in the text. This operation is performed for 10% of tokens. Synonyms for synonym replacement and random insertion are obtained from Synsets in NLTK (Bird and Loper, 2004) which are constructed based on WordNet (Miller, 1995). Synsets serve as a synonym dictionary containing groupings of synonymous words. Some words have only one Synset and some have several. In synonym replacement, if a selected word in a sentence has multiple synonyms, we randomly choose one of them, and replace all occurrences of this word in the sentence

Domain	Dataset	Label Type	Train	Dev	Test	Classes
BIO MED	CHEMPROT	relation classification	4169	2427	3469	13
	RCT	abstract sent. roles	180040	30212	30135	5
CS	ACL-ARC	citation intent	1688	114	139	6
	SciERC	relation classification	3219	455	974	7
NEWS	HYPERPARTISAN	partisanship	515	65	65	2
	AGNEWS	topic	115000	5000	7600	4
REVIEWS	HELPFULNESS	review helpfulness	115251	5000	25000	2
	IMDB	review sentiment	20000	5000	25000	2

Table 1: Statistics of datasets used in (Gururangan et al., 2020). Sources: CHEMPROT (Kringelum et al., 2016), RCT (Dernoncourt and Lee, 2017), ACL-ARC (Jurgens et al., 2018), SciERC (Luan et al., 2018), HYPERPARTISAN (Kiesel et al., 2019), AGNEWS (Zhang et al., 2015), HELPFULNESS (McAuley et al., 2015), IMDB (Maas et al., 2011). This table is taken from (Gururangan et al., 2020).

	CoLA	RTE	QNLI	STS-B	MRPC	WNLI	SST-2	MNLI (m/mm)	QQP	AX
Train	8551	2490	104743	5749	3668	635	67349	392702	363871	–
Dev	1043	277	5463	1500	408	71	872	9815/9832	40432	–
Test	1063	3000	5463	1379	1725	146	1821	9796/9847	390965	1104

Table 2: GLUE dataset statistics.

with this synonym. Random swap randomly chooses two tokens in a text and swaps their positions. This operation is performed for 10% of token pairs. Random deletion randomly removes a token with a probability of 0.1. In this SSL task, an augmented sentence is fed into a text encoder and the encoding is fed into a 4-way classification head to predict which operation was applied to generate this augmented sentence.

### 3.3 Text Encoder

We use a Transformer encoder to perform the study while noting that other text encoders are also applicable. Different from sequence-to-sequence models (Sutskever et al., 2014) that are based on recurrent neural networks (e.g., LSTM [Hochreiter and Schmidhuber, 1997], GRU [Chung et al., 2014]), which model a sequence of tokens via a recurrent manner and hence are computationally inefficient, Transformer eschews recurrent computation and instead uses self-attention which not only can capture dependency between tokens but also is amenable for parallel computation with high efficiency. Self-attention calculates the correlation among every pair of

tokens and uses these correlation scores to create “attentive” representations by taking weighted summation of tokens’ embeddings. Transformer is composed of building blocks, each consisting of a self-attention layer and a position-wise feed-forward layer. Residual connection (He et al., 2016) is applied around each of these two sub-layers, followed by layer normalization (Ba et al., 2016). Given an input sequence, an encoder—which is a stack of such building blocks—is applied to obtain a representation for each token.

## 4 Experiments

### 4.1 Datasets

We evaluated our method on the datasets used in Gururangan et al. (2020), which are from various domains. For each dataset, we follow the train/development/test split specified in Gururangan et al. (2020). Dataset statistics are summarized in Table 1.

In addition, we performed experiments on the datasets in the GLUE benchmark (Wang et al., 2018). The General Language Understanding Evaluation (GLUE) benchmark has 10 tasks, including 2 single-sentence tasks, 3 similarity and

paraphrase tasks, and 5 inference tasks. For each GLUE task, labels in development sets are publicly available and those in test sets are not released. We obtain performance on test sets by submitting inference results to GLUE evaluation server.<sup>1</sup> Table 2 shows the statistics of data split in each task.

## 4.2 Experimental Setup

### 4.2.1 Baselines

For experiments on datasets used in Gururangan et al. (2020), text encoders in all methods are initialized using pretrained RoBERTa (Liu et al., 2019a). For experiments on GLUE datasets, text encoders are initialized using pretrained BERT (Liu et al., 2019a) or pretrained RoBERTa. We compare our proposed SSL-Reg with the following baselines.

- **Unregularized RoBERTa** (Liu et al., 2019b). In this approach, the Transformer encoder is initialized with pretrained RoBERTa. Then the pretrained encoder and a classification head form a text classification model, which is then finetuned on a target classification task. Architecture of the classification model is the same as that in Liu et al. (2019b). Specifically, representation of the [CLS] special token is passed to a feedforward layer for class prediction. Non-linear activation function in the feedforward layer is tanh. During finetuning, no SSL-based regularization is used. This approach is evaluated on all datasets used in Gururangan et al. (2020) and all datasets in GLUE.
- **Unregularized BERT**. This approach is the same as unregularized RoBERTa, except that the Transformer encoder is initialized by pretrained BERT (Devlin et al., 2019a) instead of RoBERTa. This approach is evaluated on all GLUE datasets.
- **Task adaptive pretraining (TAPT)** (Gururangan et al., 2020). In this approach, given the Transformer encoder pretrained using RoBERTa or BERT on large-scale external corpora, it is further pretrained by RoBERTa or BERT on input texts in a target classification dataset (without using class labels). Then this further pretrained encoder is used to initialize the encoder in

the text classification model and is finetuned to perform classification tasks which use both input texts and their class labels. Similar to SSL-Reg, TAPT also performs SSL on texts in target classification dataset. The difference is: TAPT performs SSL task and classification task sequentially while SSL-Reg performs these two tasks jointly. TAPT is studied for all datasets in this paper.

- **Domain adaptive pretraining (DAPT)** (Gururangan et al., 2020). In this approach, given a pretrained encoder on large-scale external corpora, the encoder is further pretrained on a small-scale corpora whose domain is similar to that of texts in a target classification dataset. Then this further pretrained encoder is finetuned in a classification task. DAPT is similar to TAPT, except that TAPT performs the second stage pretraining on texts  $T$  in the classification dataset while DAPT performs the second stage pretraining on external texts whose domain is similar to that of  $T$  rather than directly on  $T$ . The external dataset is usually much larger than  $T$ .
- **TAPT+SSL-Reg**. When finetuning the classification model, SSL-Reg is applied. The rest is the same as TAPT.
- **DAPT+SSL-Reg**. When finetuning the classification model, SSL-Reg is applied. The rest is the same as DAPT.

### 4.2.2 Hyperparameter Settings

Hyperparameters were tuned on development datasets.

**Hyperparameter settings for RoBERTa on datasets used in** Gururangan et al. (2020). For a fair comparison, most of our hyperparameters are the same as those in Gururangan et al. (2020). The maximum text length was set to 512. Text encoders in all methods are initialized using pretrained RoBERTa (Liu et al., 2019a) on a large-scale external dataset. For TAPT, DAPT, TAPT+SSL-Reg, and DAPT+SSL-Reg, the second-stage pretraining on texts  $T$  in a target classification dataset or on external texts whose domain is similar to that of  $T$  is based on the pretraining approach in RoBERTa. In SSL-Reg, the SSL task is masked token prediction. SSL loss function only considers

<sup>1</sup><https://gluebenchmark.com/leaderboard>.

Task	Epoch	Learning Rate	Regularization Parameter
CoLA	10	3e-5	0.2
SST-2	3	3e-5	0.05
MRPC	5	4e-5	0.05
STS-B	10	4e-5	0.1
QQP	5	3e-5	0.2
MNLI	3	3e-5	0.1
QNLI	4	4e-5	0.5
RTE	10	3e-5	0.1
WNLI	5	5e-5	2

Table 3: Hyperparameter settings for BERT on GLUE datasets, where the SSL task is MTP.

the prediction of masked tokens and ignores the prediction of non-masked tokens. Probability for masking tokens is 0.15. If a token  $t$  is chosen to be masked, 80% of the time, we replace  $t$  with a special token [MASK]; 10% of the time, we replace  $t$  with a random word; and for the rest 10% of the time, we keep  $t$  unchanged. For the regularization parameter in SSL-Reg, we set it to 0.2 for ACL-ARC, 0.1 for SciERC, CHEMPROT, AGNEWS, RCT, HELPFULNESS, IMDB, and 0.01 for HYPERPARTISAN. For ACL-ARC, CHEMPROT, RCT, SciERC, and HYPERPARTISAN, we trained SSL-Reg for 10 epochs; for HELPFULNESS, 5 epochs; for AGNEWS, RCT and IMDB, 3 epochs. For all datasets, we used a batch size of 16 with gradient accumulation. We used the AdamW optimizer (Loshchilov and Hutter, 2017) with a warm-up proportion of 0.06, a weight decay of 0.1, and an epsilon of  $1e-6$ . In AdamW,  $\beta_1$  and  $\beta_2$  are set to 0.9 and 0.98, respectively. The maximum learning rate was  $2e-5$ .

**Hyperparameter settings for BERT on GLUE datasets.** The maximum text length was set to 128. Since external texts whose domains are similar to those of the GLUE texts are not available, we did not compare with DAPT and DAPT+SSL-Reg. For each method applied, text encoder is initialized using pretrained BERT (Devlin et al., 2019a) (with 24 layers) on a large-scale external dataset. In TAPT, the second-stage pretraining is performed using BERT. As we will show later on, TAPT does not perform well on GLUE datasets; therefore, we did not apply TAPT+SSL-Reg on these datasets further. In SSL-Reg, we studied two SSL tasks: masked token prediction (MTP) and sentence augmentation type predic-

Task	Epoch	Learning Rate	Regularization Parameter
CoLA	6	3e-5	0.4
SST-2	3	3e-5	0.8
MRPC	5	4e-5	0.05
STS-B	10	4e-5	0.05
QQP	5	3e-5	0.4
MNLI	4	3e-5	0.5
QNLI	4	4e-5	0.05
RTE	8	3e-5	0.6
WNLI	5	5e-5	0.1

Table 4: Hyperparameter settings for BERT on GLUE datasets, where the SSL task is SATP.

Task	Epoch	Learning Rate	Regularization Parameter
CoLA	10	1e-5	0.8
SST-2	3	1e-5	1.0
MRPC	10	1e-5	0.01
STS-B	10	2e-5	0.01
QQP	10	1e-5	0.1
MNLI	3	1e-5	0.1
QNLI	3	1e-5	0.1
RTE	10	2e-5	0.1
WNLI	10	2e-5	0.02

Table 5: Hyperparameter settings for RoBERTa on GLUE datasets, where the SSL task is MTP.

tion (SATP). In MTP, we randomly mask 15% of tokens in each text. Batch size was set to 32 with gradient accumulation. We use the AdamW optimizer (Loshchilov and Hutter, 2017) with a warm-up proportion of 0.1, a weight decay of 0.01, and an epsilon of  $1e-8$ . In AdamW,  $\beta_1$  and  $\beta_2$  are set to 0.9 and 0.999, respectively. Other hyperparameter settings are presented in Table 3 and Table 4.

**Hyperparameter settings for RoBERTa on GLUE datasets.** Most hyperparameter settings follow those in RoBERTa experiments performed on datasets used in Gururangan et al. (2020).

We set different learning rates and different epoch numbers for different datasets as guided by Liu et al. (2019b). In addition, we set different regularization parameters for different datasets. These hyperparameters are listed in Table 5.

Dataset	RoBERTa	DAPT	TAPT	SSL-Reg	TAPT+SSL-Reg	DAPT+SSL-Reg
CHEMPROT	81.9 <sub>1.0</sub>	84.2 <sub>0.2</sub>	82.6 <sub>0.4</sub>	83.1 <sub>0.5</sub>	83.5 <sub>0.1</sub>	<b>84.4</b> <sub>0.3</sub>
RCT	87.2 <sub>0.1</sub>	87.6 <sub>0.1</sub>	<b>87.7</b> <sub>0.1</sub>	87.4 <sub>0.1</sub>	<b>87.7</b> <sub>0.1</sub>	<b>87.7</b> <sub>0.1</sub>
ACL-ARC	63.0 <sub>5.8</sub>	75.4 <sub>2.5</sub>	67.4 <sub>1.8</sub>	69.3 <sub>4.9</sub>	68.1 <sub>2.0</sub>	<b>75.7</b> <sub>1.4</sub>
SciERC	77.3 <sub>1.9</sub>	80.8 <sub>1.5</sub>	79.3 <sub>1.5</sub>	81.4 <sub>0.8</sub>	80.4 <sub>0.6</sub>	<b>82.3</b> <sub>0.8</sub>
HYPERPARTISAN	86.6 <sub>0.9</sub>	88.2 <sub>5.9</sub>	90.4 <sub>5.2</sub>	92.3 <sub>1.4</sub>	<b>93.2</b> <sub>1.8</sub>	90.7 <sub>3.2</sub>
AGNEWS	93.9 <sub>0.2</sub>	93.9 <sub>0.2</sub>	<b>94.5</b> <sub>0.1</sub>	94.2 <sub>0.1</sub>	94.4 <sub>0.1</sub>	94.0 <sub>0.1</sub>
HELPFULNESS	65.1 <sub>3.4</sub>	66.5 <sub>1.4</sub>	68.5 <sub>1.9</sub>	69.4 <sub>0.2</sub>	<b>71.0</b> <sub>1.0</sub>	68.3 <sub>1.4</sub>
IMDB	95.0 <sub>0.2</sub>	95.4 <sub>0.1</sub>	95.5 <sub>0.1</sub>	95.7 <sub>0.1</sub>	<b>96.1</b> <sub>0.1</sub>	95.4 <sub>0.1</sub>

Table 6: Results on datasets used in Gururangan et al. (2020). For vanilla (unregularized) RoBERTa, DAPT, and TAPT, results are taken from Gururangan et al. (2020). For each method on each dataset, we run it for four times with different random seeds. Results are in  $m_s$  format, where  $m$  denotes mean and  $s$  denotes standard derivation. Following Gururangan et al. (2020), for CHEMPROT and RCT, we report micro-F1; for other datasets, we report macro-F1.

### 4.3 Results

#### 4.3.1 Results on the Datasets used in Gururangan et al. (2020)

Performance of text classification on datasets used in Gururangan et al. (2020) is reported in Table 6. Following Gururangan et al. (2020), for CHEMPROT and RCT, we report micro-F1; for other datasets, we report macro-F1. From this table, we make the following observations. First, SSL-Reg outperforms unregularized RoBERTa significantly on all datasets. We used a double-sided t-test to perform significance tests. The p-values are less than 0.01, which indicate strong statistical significance. This demonstrates the effectiveness of our proposed SSL-Reg approach in alleviating overfitting and improving generalization performance. To further confirm this, we measure the difference between F1 scores on the training set and test set in Table 7. A larger difference implies more overfitting: performing well on the training set and less well on the test set. As can be seen, the train-test difference under SSL-Reg is smaller than that under RoBERTa. SSL-Reg encourages text encoders to solve an additional task based on SSL, which reduces the risk of overfitting to the data-deficient classification task on small-sized training data. In Figure 2, we compare the training dynamics of unregularized RoBERTa and SSL-Reg (denoted by ‘‘Regularized’’). As can be seen, under a large regularization parameter  $\lambda = 1$ , our method achieves smaller

Dataset	RoBERTa	SSL-Reg
CHEMPROT	<b>13.05</b>	13.57
ACL-ARC	28.67	<b>25.24</b>
SciERC	19.51	<b>18.23</b>
HYPERPARTISAN	7.44	<b>5.64</b>

Table 7: Difference between F1 score on training set and F1 score on test set with or without SSL-Reg (MTP). **Bold** denotes a smaller difference, which means overfitting is less severe.

differences between training accuracy and validation accuracy than unregularized RoBERTa; our method also achieves smaller differences between training accuracy and test accuracy than unregularized RoBERTa. These results show that our proposed SSL-Reg indeed acts as a regularizer which reduces the gap between performances on training set and validation/test set. Besides, when increasing  $\lambda$  from 0.1 to 1, the training accuracy of SSL-Reg decreases considerably. This also indicates that SSL-Reg acts as a regularizer which penalizes training performance.

Second, on 6 out of the 8 datasets, SSL-Reg performs better than TAPT. On the other two datasets, SSL-Reg is on par with TAPT. This shows that SSL-Reg is more effective than TAPT. SSL-Reg and TAPT both leverage input texts in classification datasets for self-supervised learning. The difference is: TAPT uses these texts to pretrain the encoder while SSL-Reg uses these texts to regularize the encoder during finetuning. In SSL-Reg, the encoder is learned to perform

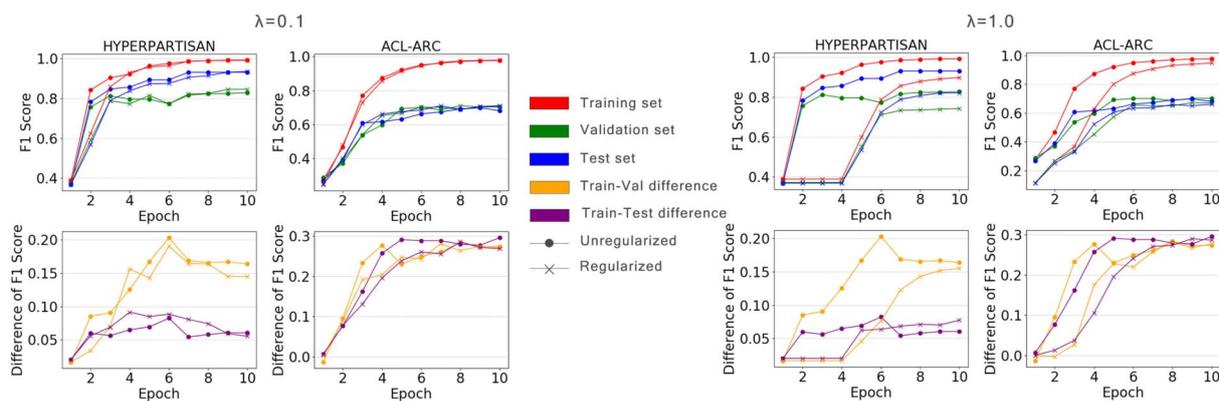


Figure 2: Training dynamics of unregularized RoBERTa and SSL-Reg (denoted by ‘‘Regularized’’) on HYPERPARTISAN and ACL-ARC. In SSL-Reg, we experimented with two values of the regularization parameter  $\lambda$ : 0.1 and 1.

classification tasks and SSL tasks simultaneously. Thus the encoder is not completely biased to classification tasks. In TAPT, the encoder is first learned by performing SSL tasks, then finetuned by performing classification tasks. There is a risk that after finetuning, the encoder is largely biased to classification tasks on small-sized training data, which leads to overfitting.

Third, on 5 out of the 8 datasets, SSL-Reg performs better than DAPT, although DAPT leverages additional external data. The reasons are two-fold: 1) similar to TAPT, DAPT performs SSL task first and then classification task separately; as a result, the encoder may be eventually biased to classification task on small-sized training data; 2) external data used in DAPT still has a domain shift with target dataset; this domain shift may render the text encoder pretrained on external data not suitable for target task. To verify this, we measure the domain similarity between external texts and target texts by calculating cosine similarity between the BERT embeddings of these texts. The similarity score is 0.14. As a reference, the similarity score between texts in the target dataset is 0.27. This shows that there is indeed a domain difference between external texts and target texts.

Fourth, on 6 out of 8 datasets, TAPT+SSL-Reg performs better than TAPT. On the other two datasets, TAPT+SSL-Reg is on par with TAPT. This further demonstrates the effectiveness of SSL-Reg.

Fifth, on all eight datasets, DAPT+SSL-Reg performs better than DAPT. This again shows that SSL-Reg is effective.

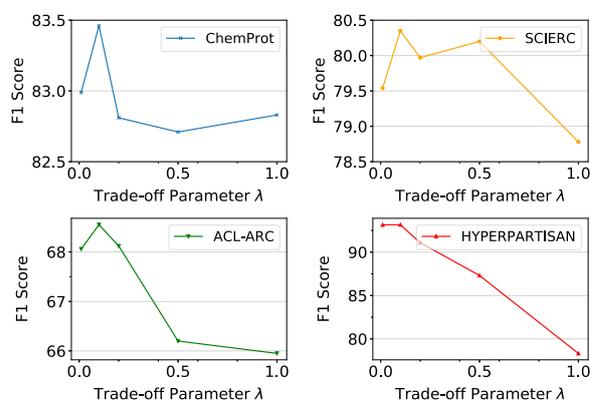


Figure 3: How regularization parameter in SSL-Reg affects text classification F1 score.

Sixth, on 6 out of 8 datasets, TAPT+SSL-Reg performs better than SSL-Reg, indicating that it is beneficial to use both TAPT and SSL-Reg: first use the target texts to pretrain the encoder based on SSL, then apply SSL-based regularizer on these target texts during finetuning.

Seventh, DAPT+SSL-Reg performs better than SSL-Reg on 4 datasets, but worse on the other 4 datasets, indicating that with SSL-Reg used, DAPT is not necessarily useful.

Eighth, on smaller datasets, improvement achieved by SSL-Reg over baselines is larger. For example, on HYPERPARTISAN which has only about 500 training examples, improvement of SSL-Reg over RoBERTa is 5.7% (absolute percentage). Relative improvement is 6.6%. As another example, on ACL-ARC which has about 1700 training examples, improvement of SSL-Reg over RoBERTa is 6.3% (absolute percentage).

	CoLA (Matthew Corr.)	SST-2 (Accuracy)	RTE (Accuracy)	QNLI (Accuracy)	MRPC (Accuracy/F1)
<i>The median result</i>					
BERT, Lan et al. 2019	60.6	93.2	70.4	92.3	<b>88.0/-</b>
BERT, our run	62.1	93.1	74.0	92.1	86.8/90.8
TAPT	61.2	93.1	74.0	92.0	85.3/89.8
SSL-Reg (SATP)	63.7	<b>93.9</b>	<b>74.7</b>	92.3	86.5/90.3
SSL-Reg (MTP)	<b>63.8</b>	93.8	<b>74.7</b>	<b>92.6</b>	87.3/90.9
<i>The best result</i>					
BERT, our run	63.9	93.3	75.8	92.5	<b>89.5/92.6</b>
TAPT	62.0	93.9	76.2	92.4	86.5/90.7
SSL-Reg (SATP)	65.3	94.6	<b>78.0</b>	92.8	88.5/91.9
SSL-Reg (MTP)	<b>66.3</b>	<b>94.7</b>	<b>78.0</b>	<b>93.1</b>	<b>89.5/92.4</b>

Table 8: Results of BERT-based experiments on GLUE development sets, where results on MNLI and QQP are the median of five runs and results on other datasets are the median of nine runs. The size of MNLI and QQP is very large, taking a long time to train on. Therefore, we reduced the number of runs. Because we used a different optimization method to re-implement BERT, our median performance is not the same as that reported in Lan et al. (2019).

Relative improvement is 10%. In contrast, on large datasets such as RCT which contains about 180000 training examples, improvement of SSL-Reg over RoBERTa is 0.2% (absolute percentage). Relative improvement is 0.2%. On another large dataset AGNEWS which contains 115000 training examples, improvement of SSL-Reg over RoBERTa is 0.3% (absolute percentage). Relative improvement is 0.3%. The reason that SSL-Reg achieves better improvement on smaller datasets is that smaller datasets are more likely to lead to overfitting and SSL-Reg is more needed to alleviate this overfitting.

Figure 3 shows how classification F1 score varies as we increase regularization parameter  $\lambda$  from 0.01 to 1.0 in SSL-Reg. As can be seen, starting from 0.01, when the regularizer parameter is increasing, F1 score increases. This is because a larger  $\lambda$  imposes a stronger regularization effect, which helps to reduce overfitting. However, if  $\lambda$  becomes too large, F1 score drops. This is because the regularization effect is too strong, which dominates classification loss. Among these 4 datasets, F1 score drops dramatically on HYPERPARTISAN as  $\lambda$  increases. This is probably because this dataset contains very long sequences. This makes MTP on this dataset more difficult and therefore yields an excessively strong regularization outcome that hurts classification performance. Compared with HYPERPARTISAN, F1 score is less

sensitive on other datasets because their sequence lengths are relatively smaller.

### 4.3.2 Results on the GLUE Benchmark

Table 8 and Table 9 show results of BERT-based experiments on development sets of GLUE. As mentioned in (Devlin et al., 2019b), for the 24-layer version of BERT, finetuning is sometimes unstable on small datasets, so we run each method several times and report the median and best performance. Table 10 shows the best performance on test sets. Following Wang et al. (2018), we report Matthew correlation on CoLA, Pearson correlation and Spearman correlation on STS-B, accuracy and F1 on MRPC and QQP. For the rest of the datasets, we report accuracy. From these tables, we make the following observations. First, SSL-Reg methods including SSL-Reg-SATP and SSL-Reg-MTP outperform unregularized BERT (our run) on most datasets: 1) on test sets, SSL-Reg-SATP performs better than BERT on 7 out of 10 datasets and SSL-Reg-MTP performs better than BERT on 9 out of 10 datasets; 2) in terms of median results on development sets, SSL-Reg-SATP performs better than BERT (our run) on 7 out of 9 datasets and SSL-Reg-MTP performs better than BERT (our run) on 8 out of 9 datasets; 3) in terms of best results on development sets, SSL-Reg-SATP performs better than BERT (our run) on 8 out of 9 datasets and SSL-Reg-MTP performs better than

	MNLI-m/mm (Accuracy)	QQP (Accuracy/F1)	STS-B (Pearson Corr./Spearman Corr.)	WNLI (Accuracy)
<i>The median result</i>				
BERT, Lan et al. 2019	86.6/–	91.3/–	90.0/–	–
BERT, our run	86.2/86.0	91.3/88.3	90.4/90.0	56.3
TAPT	85.6/85.5	91.5/88.7	90.6/90.2	53.5
SSL-Reg (SATP)	86.2/86.2	91.6/88.8	<b>90.7/90.4</b>	56.3
SSL-Reg (MTP)	<b>86.6/86.6</b>	<b>91.8/89.0</b>	90.7/90.3	56.3
<i>The best result</i>				
BERT, our run	86.4/86.3	91.4/88.4	90.9/90.5	56.3
TAPT	85.7/85.7	91.7/89.0	90.8/90.4	56.3
SSL-Reg (SATP)	86.4/86.5	91.8/88.9	<b>91.1/90.8</b>	<b>59.2</b>
SSL-Reg (MTP)	<b>86.9/86.9</b>	<b>91.9/89.1</b>	<b>91.1/90.8</b>	57.7

Table 9: Continuation of Table 8.

	BERT	TAPT	SSL-Reg (SATP)	SSL-Reg (MTP)
CoLA (Matthew Corr.)	60.5	61.3	<b>63.0</b>	61.2
SST-2 (Accuracy)	94.9	94.4	95.1	<b>95.2</b>
RTE (Accuracy)	70.1	70.3	71.2	<b>72.7</b>
QNLI (Accuracy)	92.7	92.4	92.5	<b>93.2</b>
MRPC (Accuracy/F1)	85.4/89.3	85.9/89.5	85.3/89.3	<b>86.1/89.8</b>
MNLI-m/mm (Accuracy)	<b>86.7/85.9</b>	85.7/84.4	86.2/85.4	86.6/ <b>86.1</b>
QQP (Accuracy/F1)	89.3/72.1	89.3/71.6	89.6/72.2	<b>89.7/72.5</b>
STS-B (Pearson Corr./Spearman Corr.)	87.6/86.5	<b>88.4/87.3</b>	88.3/ <b>87.5</b>	88.1/87.2
WNLI (Accuracy)	65.1	65.8	65.8	<b>66.4</b>
AX(Matthew Corr.)	39.6	39.3	40.2	<b>40.3</b>
<b>Average</b>	80.5	80.6	81.0	<b>81.3</b>

Table 10: Results of BERT-based experiments on GLUE test sets, which are scored by the GLUE evaluation server (<https://gluebenchmark.com/leaderboard>). Models evaluated on AX are trained on the training dataset of MNLI.

BERT (our run) on 8 out of 9 datasets. This further demonstrates the effectiveness of SSL-Reg in improving generalization performance.

Second, on 7 out of 10 test sets, SSL-Reg-SATP outperforms TAPT; on 8 out of 10 test sets, SSL-Reg-MTP outperforms TAPT. On most development datasets, SSL-Reg-SATP and SSL-Reg-MTP outperform TAPT. The only exception is: on QQP development set, the best F1 of TAPT is slightly better than that of SSL-Reg-SATP. This further demonstrates that performing SSL-based regularization on target texts is more effective than using them for pretraining.

Third, overall, SSL-Reg-MTP performs better than SSL-Reg-SATP. For example, on 8 out of 10 test datasets, SSL-Reg-MTP performs better than

SSL-Reg-SATP. MTP works better than SATP probably because it is a more challenging self-supervised learning task that encourages encoders to learn more powerful representations.

Fourth, improvement of SSL-Reg methods over BERT is more prominent on smaller training datasets, such as CoLA and RTE. This may be because smaller training datasets are more likely to lead to overfitting where the advantage of SSL-Reg in alleviating overfitting can be better played.

Tables 11 and 12 show results of RoBERTa-based experiments on development sets of GLUE. From these two tables, we make observations that are similar to those in Table 8 and Table 9. In terms of median results, SSL-Reg (MTP) performs better than unregularized RoBERTa (our

	CoLA (Matthew Corr.)	SST-2 (Accuracy)	RTE (Accuracy)	QNLI (Accuracy)	MRPC (Accuracy/F1)
<i>The median result</i>					
RoBERTa, Liu et al. 2019	68.0	<b>96.4</b>	<b>86.6</b>	94.7	<b>90.9/-</b>
RoBERTa, our run	68.7	96.1	84.8	94.6	89.5/92.3
SSL-Reg (MTP)	<b>69.2</b>	96.3	85.2	<b>94.9</b>	90.0/ <b>92.7</b>
<i>The best result</i>					
RoBERTa, our run	69.2	96.7	86.6	94.7	90.4/93.1
SSL-Reg (MTP)	<b>70.2</b>	96.7	86.6	<b>95.2</b>	<b>91.4/93.8</b>

Table 11: Results of RoBERTa-based experiments on GLUE development sets, where the median results are the median of five runs. Because we used a different optimization method to re-implement RoBERTa, our median performance is not the same as that reported in (Liu et al., 2019b).

	MNLI-m/mm (Accuracy)	QQP (Accuracy)	STS-B (Pearson Corr./ Spearman Corr.)	WNLI (Accuracy)
<i>The median result</i>				
RoBERTa, Liu et al. 2019	90.2/90.2	92.2	92.4/-	-
RoBERTa, our run	90.5/90.5	91.6	92.0/92.0	56.3
SSL-Reg (MTP)	<b>90.7/90.7</b>	91.6	92.0/92.0	<b>62.0</b>
<i>The best result</i>				
RoBERTa, our run	90.7/90.5	91.7	92.3/92.2	60.6
SSL-Reg (MTP)	90.7/90.5	<b>91.8</b>	92.3/92.2	<b>66.2</b>

Table 12: Continuation of Table 11.

	CoLA	SST-2	RTE	QNLI	MRPC	STS-B
SR+RD+RI+RS	<b>63.6</b>	<b>94.0</b>	<b>74.8</b>	<b>92.2</b>	86.8/90.6	<b>90.6/90.3</b>
SR+RD+RI	63.4	93.8	72.8	92.1	86.9/90.8	90.6/90.2
SR+RD	61.6	93.6	72.5	<b>92.2</b>	<b>87.2/91.0</b>	<b>90.6/90.3</b>

Table 13: Ablation study on sentence augmentation types in SSL-Reg (SATP), where SR, RD, RI, and RS denotes synonym replacement, random deletion, random insertion, and random swap respectively. Results are averaged over 5 runs with different random initialization.

run) on 7 out of 9 datasets and achieves the same performance as RoBERTa (our run) on the rest 2 datasets. In terms of best results, SSL-Reg (MTP) performs better than RoBERTa (our run) on 5 out of 9 datasets and achieves the same performance as RoBERTa (our run) on the rest 4 datasets. This further demonstrates the effectiveness of our proposed SSL-Reg approach which uses an MTP-based self-supervised task to regularize the finetuning of RoBERTa.

In SSL-Reg (SATP), we perform an ablation study on different types of sentence augmenta-

tion. Results are shown in Table 13, where SR, RD, RI, and RS denote synonym replacement, random deletion, random insertion, and random swap, respectively. SR+RD+RI+RS means that we apply these four types of operations to augment sentences; given an augmented sentence  $a$ , we predict which of the four types of operations was applied to an original sentence to create  $a$ . SR+RD+RI+RS and SR+RD hold similar meanings. From this table, we make the following observations.

First, as the number of augmentation types increases from 2 (SR+RD) to 3 (SR+RD+RI) then

to 4 (SR+RD+RI+RS), the performance increases in general. This shows that it is beneficial to have more augmentation types in SATP. The reason is that more types make the SATP task more challenging and solving a more challenging self-supervised learning task can enforce sentence encoders to learn more powerful representations.

Second, SR+RD+RI+RS outperforms SR+RD+RI on 5 out of 6 datasets. This demonstrates that leveraging random swap (RS) for SATP can learn more effective representations of sentences. The reason is: SR, RD, and RI change the collection of tokens in a sentence via synonym replacement, random deletion, and random insertion; RS does not change the collection of tokens, but changes the order of tokens; therefore, RS is complementary to the other three operations; adding RS can bring in additional benefits that are complementary to those of SR, RD, and RI.

Third, SR+RD+RI performs much better than SR+RD on CoLA and is on par with SR+RD on the rest five datasets. This shows that adding RI to SR+RD is beneficial. Unlike synonym replacement (SR) and random deletion (RD) which do not increase the number of tokens in a sentence, RI increases token number. Therefore, RI is complementary to SR and RD and can bring in additional benefits.

## 5 Conclusions and Future Work

In this paper, we propose to use self-supervised learning to alleviate overfitting in text classification problems. We propose SSL-Reg, which is a regularizer based on SSL and a text encoder is trained to simultaneously minimize classification loss and regularization loss. We demonstrate the effectiveness of our methods on 17 text classification datasets.

For future work, we will use other self-supervised learning tasks to perform regularization, such as contrastive learning, which predicts whether two augmented sentences stem from the same original sentence.

## References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Philip Bachman, R. Devon Hjelm, and William Buchwalter. 2019. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15509–15519.
- Steven Bird and Edward Loper. 2004. NLTK: The natural language toolkit. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics. <https://doi.org/10.3115/1219044.1219075>
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Franck Dernoncourt and Ji Young Lee. 2017. Pubmed 200k RCT: A dataset for sequential sentence classification in medical abstracts. In *IJCNLP*. <https://doi.org/10.18653/v1/E17-2110>
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019a. BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL-HLT*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019b. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Hongchao Fang, Sicheng Wang, Meng Zhou, Jiayuan Ding, and Pengtao Xie. 2020. Cert: Contrastive self-supervised learning for language understanding. *arXiv e-prints arXiv:2005.12766*. <https://doi.org/10.36227/techrxiv.12308378.v1>
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. 2018. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't stop pretraining: Adapt language models to

- domains and tasks. In *Proceedings of ACL*. <https://doi.org/10.18653/v1/2020.acl-main.740>
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2019. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778.
- Xuehai He, Zhuo Cai, Wenlan Wei, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. 2020a. Pathological visual question answering. *arXiv preprint arXiv:2010.12435*.
- Xuehai He, Xingyi Yang, Shanghang Zhang, Jinyu Zhao, Yichen Zhang, Eric Xing, and Pengtao Xie. 2020b. Sample-efficient deep learning for covid-19 diagnosis based on CT scans. *medRxiv*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>, PubMed: 9377276
- J. Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *ACL*. <https://doi.org/10.18653/v1/P18-1031>
- David Jurgens, Srijan Kumar, Raine Hoover, Daniel A. McFarland, and Dan Jurafsky. 2018. Measuring the evolution of a scientific field through citation frames. *TACL*. <https://doi.org/10.1162/tacl.a.00028>
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 655–665, Baltimore, Maryland. Association for Computational Linguistics. <https://doi.org/10.3115/v1/P14-1062>
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*.
- Johannes Kiesel, Maria Mestre, Rishabh Shukla, Emmanuel Vincent, Payam Adineh, David Corney, Benno Stein, and Martin Potthast. 2019. SemEval-2019 Task 4: Hyperpartisan news detection. In *SemEval*. <https://doi.org/10.18653/v1/S19-2145>
- Tassilo Klein and Moin Nabi. 2020. Contrastive self-supervised learning for commonsense reasoning. *arXiv preprint arXiv:2005.00669*. <https://doi.org/10.18653/v1/2020.acl-main.671>
- V. Korde and C. Mahender. 2012. Text classification and classifiers: A survey. *International Journal of Artificial Intelligence & Applications*, 3:85–99. <https://doi.org/10.5121/ijaiia.2012.3208>
- Jens Kringelum, Sonny Kim Kjærulff, Søren Brunak, Ole Lund, Tudor I. Oprea, and Olivier Taboureau. 2016. ChemProt-3.0: A global chemical biology diseases mapping. In *Database*. <https://doi.org/10.1093/database/bav123>, PubMed: 26876982
- Siwei Lai, L. Xu, Kang Liu, and Jun Zhao. 2015. Recurrent convolutional neural networks for text classification. In *AAAI*.
- 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, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Xueting Li, Sifei Liu, Shalini De Mello, Xiaolong Wang, Jan Kautz, and Ming-Hsuan Yang. 2019. Joint-task self-supervised learning for temporal correspondence. In *Advances in Neural Information Processing Systems*, pages 317–327.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. Recurrent neural network for text

- classification with multi-task learning. *arXiv preprint arXiv:1605.05101*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019a. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv:1907.11692*.
- I. Loshchilov and F. Hutter. 2017. Fixing weight decay regularization in Adam. *ArXiv*, abs/1711.05101.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In *EMNLP*. <https://doi.org/10.18653/v1/D18-1360>
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *ACL*.
- Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *ACM SIGIR*. <https://doi.org/10.1145/2766462.2767755>
- G. Miller. 1995. Wordnet: A lexical database for English. *Communications of ACM*, 38:39–41.
- Shervin Minaee, Nal Kalchbrenner, Erik Cambria, Narjes Nikzad, Meysam Chenaghlu, and Jianfeng Gao. 2020. Deep learning based text classification: A comprehensive review. *arXiv preprint arXiv:2004.03705*.
- T. Nathan Mundhenk, Daniel Ho, and Barry Y. Chen. 2018. Improvements to context based self-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9339–9348. <https://doi.org/10.1109/CVPR.2018.00973>
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544. <https://doi.org/10.1109/CVPR.2016.278>
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. <https://www.cs.ubc.ca/~amuham01/LING530/papers/radford2018improving.pdf>
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Sam T. Roweis and Lawrence K. Saul. 2000. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500): 2323–2326. <https://doi.org/10.1126/science.290.5500.2323>, PubMed: 11125150
- Aravind Srinivas, Michael Laskin, and Pieter Abbeel. 2020. Curl: Contrastive unsupervised representations for reinforcement learning. *arXiv preprint arXiv:2004.04136*.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Hao Tian, Hua Wu, and Haifeng Wang. 2019. Ernie 2.0: A continual pre-training framework for language understanding. *arXiv preprint arXiv:1907.12412*. <https://doi.org/10.1609/aaai.v34i05.6428>
- Yu Sun, Xiaolong Wang, Liu Zhuang, John Miller, Moritz Hardt, and Alexei A. Efros. 2020. Test-time training with self-supervision for generalization under distribution shifts. In *ICML*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, pages 3104–3112.

- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1556–1566, Beijing, China. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*. <https://doi.org/10.18653/v1/W18-5446>
- Jin Wang, Zhongyuan Wang, D. Zhang, and Jun Yan. 2017. Combining knowledge with deep convolutional neural networks for short text classification. In *IJCAI*. <https://doi.org/10.24963/ijcai.2017/406>
- Xiaolong Wang, Allan Jabri, and Alexei A. Efros. 2019b. Learning correspondence from the cycle-consistency of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2566–2576. <https://doi.org/10.1109/CVPR.2019.00267>
- Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. 2019a. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6629–6638. <https://doi.org/10.1109/CVPR.2019.00679>
- Jason Wei and Kai Zou. 2019. EDA: Easy data augmentation techniques for boosting performance on text classification tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6383–6389, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1670>
- Qingyang Wu, Lei Li, Hao Zhou, Ying Zeng, and Zhou Yu. 2019. Importance-aware learning for neural headline editing. *arXiv preprint arXiv:1912.01114*.
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742.
- Xingyi Yang, Xuehai He, Yuxiao Liang, Yue Yang, Shanghang Zhang, and Pengtao Xie. 2020. Transfer learning or self-supervised learning? a tale of two pretraining paradigms. *arXiv preprint arXiv:2007.04234*. <https://doi.org/10.36227/techrxiv.12502298>
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R. Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.
- Jiaqi Zeng and Pengtao Xie. 2021. Contrastive self-supervised learning for graph classification. *AAAI*.
- Richard Zhang, Phillip Isola, and Alexei A. Efros. 2016. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer. [https://doi.org/10.1007/978-3-319-46487-9\\_40](https://doi.org/10.1007/978-3-319-46487-9_40)
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *NeurIPS*.
- Chunting Zhou, Chonglin Sun, Zhiyuan Liu, and F. C. M. Lau. 2015. A c-lstm neural network for text classification. *ArXiv*, abs/1511.08630.