

SentiLARE: Sentiment-Aware Language Representation Learning with Linguistic Knowledge

Pei Ke*, Haozhe Ji*, Siyang Liu, Xiaoyan Zhu, Minlie Huang†

Department of Computer Science and Technology, Institute for Artificial Intelligence,
State Key Lab of Intelligent Technology and Systems,

Beijing National Research Center for Information Science and Technology,

Tsinghua University, Beijing 100084, China

kepei1106@outlook.com, {jhz20, siyang-118}@mails.tsinghua.edu.cn
{zxy-dcs, aihuang}@tsinghua.edu.cn

Abstract

Most of the existing pre-trained language representation models neglect to consider the linguistic knowledge of texts, which can promote language understanding in NLP tasks. To benefit the downstream tasks in sentiment analysis, we propose a novel language representation model called SentiLARE, which introduces word-level linguistic knowledge including part-of-speech tag and sentiment polarity (inferred from SentiWordNet) into pre-trained models. We first propose a context-aware sentiment attention mechanism to acquire the sentiment polarity of each word with its part-of-speech tag by querying SentiWordNet. Then, we devise a new pre-training task called label-aware masked language model to construct knowledge-aware language representation. Experiments show that SentiLARE obtains new state-of-the-art performance on a variety of sentiment analysis tasks¹.

1 Introduction

Recently, pre-trained language representation models such as GPT (Radford et al., 2018, 2019), ELMo (Peters et al., 2018), and BERT (Devlin et al., 2019) have achieved promising results in NLP tasks, including sentiment analysis (Xu et al., 2019, 2020; Yin et al., 2020). These models capture contextual information from large-scale corpora via well-designed pre-training tasks. The literature has commonly reported that pre-trained models can be used as effective feature extractors and achieve state-of-the-art performance on various downstream tasks (Wang et al., 2019a).

Despite the great success of pre-trained models, existing pre-training tasks like masked language model and next sentence prediction (Devlin et al.,

2019) neglect to consider the linguistic knowledge. Such knowledge is important for some NLP tasks, particularly for sentiment analysis. For instance, existing work has shown that linguistic knowledge including part-of-speech tag (Qian et al., 2015; Huang et al., 2017) and word-level sentiment polarity (Qian et al., 2017) is closely related to the sentiment of longer texts. We argue that pre-trained models enriched with the linguistic knowledge of words will facilitate the understanding of the sentiment of the whole texts, thereby resulting in better performance on sentiment analysis.

There are two major challenges to construct knowledge-aware pre-trained language representation models which can promote the downstream tasks in sentiment analysis: 1) **Knowledge acquisition across different contexts**. Most of the existing work has adopted static sentiment lexicons as linguistic resource (Qian et al., 2017; Chen et al., 2019), and equipped each word with a fixed sentiment polarity across different contexts. However, the same word may play different sentiment roles in different contexts due to the variety of part-of-speech tags and word senses. 2) **Knowledge integration into pre-trained models**. Since the introduced word-level linguistic knowledge can only reflect the local sentiment role played by each word, it is important to deeply integrate knowledge into pre-trained models to construct sentence-level language representation, which can derive the global sentiment label of a whole sentence from local information. How to build the connection between sentence-level language representation and word-level linguistic knowledge is underexplored.

In this paper, we propose a novel pre-trained language representation model called SentiLARE to deal with these challenges. First, to acquire the linguistic knowledge of each word, we label the word with its part-of-speech tag, and obtain the sentiment polarity via a context-aware sentiment

* Equal contribution

† Corresponding author

¹The data, codes, and model parameters are available at <https://github.com/thu-coai/SentiLARE>.

attention mechanism over all the matched senses in SentiWordNet (Baccianella et al., 2010). Then, to incorporate linguistic knowledge into pre-trained models, we devise a novel pre-training task called label-aware masked language model. This task involves two sub-tasks: 1) predicting the word, part-of-speech tag, and sentiment polarity at masked positions given the sentence-level sentiment label; 2) predicting the sentence-level label, the masked word and its linguistic knowledge including part-of-speech tag and sentiment polarity simultaneously. We call the first sub-task *early fusion* since the sentiment labels are beforehand integrated as input embeddings, whereas in the second sub-task, the labels are used as *late supervision* to the model in the output layer. These two sub-tasks are expected to establish the connection between sentence-level representation and word-level linguistic knowledge, which can benefit downstream tasks in sentiment analysis. Our contributions are in three folds:

- We investigate the effectiveness of incorporating linguistic knowledge into pre-trained language representation models, and we reveal that injecting such knowledge via pre-training tasks can benefit various downstream tasks in sentiment analysis.
- We propose a novel pre-trained language representation model called SentiLARE. This model derives a context-aware sentiment polarity for each word using SentiWordNet, and adopts a pre-training task named label-aware masked language model to construct sentiment-aware language representations.
- We conduct extensive experiments on sentence-level and aspect-level sentiment analysis (including extraction and classification). Results show that SentiLARE obtains new state-of-the-art performance on a variety of sentiment analysis tasks.

2 Related Work

General Pre-trained Language Models

Recently, pre-trained language representation models including ELMo (Peters et al., 2018), GPT (Radford et al., 2018, 2019), and BERT (Devlin et al., 2019) become prevalent. These models use LSTM (Hochreiter and Schmidhuber, 1997) or Transformer (Vaswani et al., 2017) as the encoder to acquire contextual language representation, and

explore various pre-training tasks including masked language model and next sentence prediction (Devlin et al., 2019).

Thanks to the great success of BERT on various NLP tasks, many variants of BERT have been proposed, which mainly fall into four aspects: 1) Knowledge enhancement: ERNIE-Tsinghua (Zhang et al., 2019) / KnowBERT (Peters et al., 2019) explicitly introduces knowledge graph / knowledge base to BERT, while ERNIE-Baidu (Sun et al., 2019b) designs entity-specific masking strategies during pre-training. 2) Transferability: TransBERT (Li et al., 2019) conducts supervised post-training on the pre-trained BERT with transfer tasks to get a better initialization for target tasks. 3) Hyper-parameters: RoBERTa (Liu et al., 2019) measures the impact of key hyper-parameters to improve the under-trained BERT. 4) Pre-training tasks: SpanBERT (Joshi et al., 2020) masks consecutive spans randomly instead of individual tokens, while XLNet (Yang et al., 2019) designs a training objective combining both reconstruction and autoregressive language modeling.

Pre-trained Models for Sentiment Analysis

Another line of work aims to build task-specific pre-trained models via post-training on the task data (Gururangan et al., 2020). For sentiment analysis, BERT-PT (Xu et al., 2019) conducts post-training on the corpora which belong to the same domain of the downstream tasks to benefit aspect-level sentiment analysis. DomBERT (Xu et al., 2020) augments the training samples from relevant domains during the pre-training phase to enhance the performance on the aspect-level sentiment analysis of target domains. SentiBERT (Yin et al., 2020) devises a two-level attention mechanism on top of the BERT representation to capture phrase-level compositional semantics.

Compared with the existing work on pre-trained models for sentiment analysis, our work integrates sentiment-related linguistic knowledge from SentiWordNet (Baccianella et al., 2010) into pre-trained models to construct knowledge-aware language representation, which can benefit a wide range of downstream tasks in sentiment analysis.

Linguistic Knowledge for Sentiment Analysis

Linguistic knowledge such as part of speech and word-level sentiment polarity is commonly used as external features in sentiment analysis. Part of speech is shown to facilitate the parsing of the syntactic structure of texts (Socher et al., 2013). It

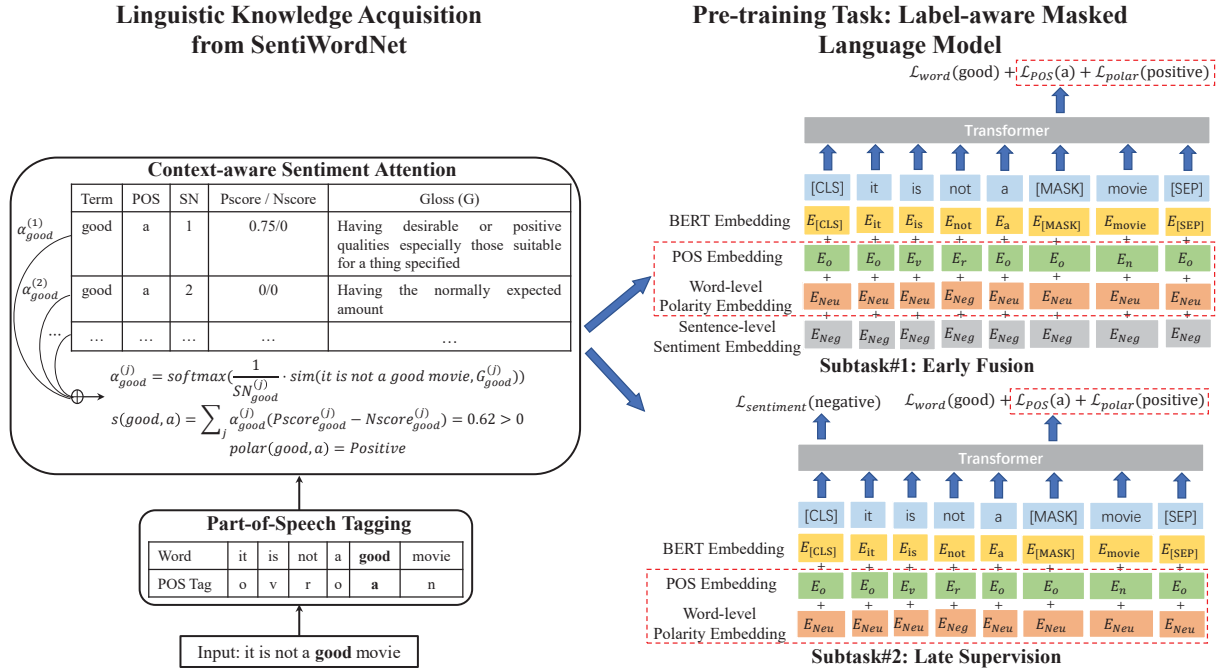


Figure 1: Overview of SentiLARE. This model first labels each word with its part-of-speech tag, and then uses the word and tag to match the corresponding senses in SentiWordNet. The sentiment polarity of each word is obtained by weighting the matched senses with context-aware sentiment attention. During pre-training, the model is trained based on label-aware masked language model including *early fusion* and *late supervision*. Red dotted boxes denote that the linguistic knowledge is used in input embedding or pre-training loss function.

can also be incorporated into all layers of RNN as tag embeddings (Qian et al., 2015). Huang et al. (2017) shows that part of speech can help to learn sentiment-favorable representations.

Word-level sentiment polarity is mostly derived from sentiment lexicons (Hu and Liu, 2004; Wilson et al., 2005). Guerini et al. (2013) obtains the prior sentiment polarity by weighting the sentiment scores over all the senses of a word in SentiWordNet (Esuli and Sebastiani, 2006; Baccianella et al., 2010). Teng et al. (2016) proposes a context-aware lexicon-based weighted sum model, which weights the prior sentiment scores of sentiment words to derive the sentiment label of the whole sentence. Qian et al. (2017) models the linguistic role of sentiment, negation and intensity words via linguistic regularizers in the training objective of LSTM.

3 Model

3.1 Task Definition and Model Overview

Our task is defined as follows: given a text sequence $X = (x_1, x_2, \dots, x_n)$ of length n , our goal is to acquire the representation of the whole sequence $H = (h_1, h_2, \dots, h_n)^T \in \mathbb{R}^{n \times d}$ that captures the contextual information and the linguistic knowledge simultaneously, where d indicates

the dimension of representation vectors.

Figure 1 shows the overview of our model, which consists of two steps: 1) Acquiring the part-of-speech tag and the sentiment polarity for each word; 2) Conducting pre-training via label-aware masked language model, which contains two pre-training sub-tasks, i.e., early fusion and late supervision. Compared with existing BERT-style pre-trained models, our model enriches the input sequence with its linguistic knowledge including part-of-speech tag and sentiment polarity, and utilizes label-aware masked language model to capture the relationship between sentence-level language representation and word-level linguistic knowledge.

3.2 Linguistic Knowledge Acquisition

This module obtains the part-of-speech tag and the sentiment polarity for each word. The input of this module is a text sequence $X = (x_1, x_2, \dots, x_n)$, where $x_i (1 \leq i \leq n)$ indicates a word in the vocabulary. First, our model acquires the part-of-speech tag pos_i of each word x_i via Stanford Log-Linear Part-of-Speech Tagger². For simplicity, we only consider five POS tags including verb (v), noun (n),

²<https://nlp.stanford.edu/software/tagger.html>

adjective (*a*), adverb (*r*), and others (*o*).

Then, we acquire the word-level sentiment polarity $polar_i$ from SentiWordNet for each pair (x_i, pos_i) . In SentiWordNet, we can find m different senses for the pair (x_i, pos_i) , each of which contains a sense number, a positive / negative score, and a gloss $(SN_i^{(j)}, Pscore_i^{(j)}, Nscore_i^{(j)}, G_i^{(j)})$, $1 \leq j \leq m$, where SN indicates the rank of different senses, $Pscore/Nscore$ is the positive / negative score assigned by SentiWordNet, and G denotes the definition of each sense. Inspired by the existing work on inferring word-level prior polarity from SentiWordNet (Guerini et al., 2013) and unsupervised word sense disambiguation (Basile et al., 2014), we propose a context-aware attention mechanism which simultaneously considers the sense rank and the context-gloss similarity to determine the attention weight of each sense:

$$\alpha_i^{(j)} = softmax\left(\frac{1}{SN_i^{(j)}} \cdot sim(X, G_i^{(j)})\right)$$

where $\frac{1}{SN_i^{(j)}}$ approximates the impact of sense frequency because a smaller sense rank indicates more frequent use of this sense in natural language (Guerini et al., 2013), and $sim(X, G_i^{(j)})$ denotes the textual similarity between the context and the gloss of each sense, which is commonly used as an important feature in unsupervised word sense disambiguation (Basile et al., 2014). To calculate the similarity between X and $G_i^{(j)}$, we encode them with Sentence-BERT (SBERT) (Reimers and Gurevych, 2019) which achieves the state-of-the-art performance on semantic textual similarity tasks, and obtain the cosine similarity between the vectors:

$$sim(X, G_i^{(j)}) = cos(SBERT(X), SBERT(G_i^{(j)}))$$

Once we obtain the attention weight of each sense, we can calculate the sentiment score of each pair (x_i, pos_i) by simply weighting the scores of all the senses:

$$s(x_i, pos_i) = \sum_{j=1}^m \alpha_i^{(j)} (Pscore_i^{(j)} - Nscore_i^{(j)})$$

Finally, the word-level sentiment polarity $polar_i$ for the pair (x_i, pos_i) can be assigned with *Positive/Negative/Neutral* when $s(x_i, pos_i)$ is positive / negative / zero, respectively. Note that if we cannot find any sense for (x_i, pos_i) in SentiWordNet, $polar_i$ is assigned with *Neutral*.

3.3 Pre-training Task

Given the knowledge enhanced text sequence $X_k = \{(x_i, pos_i, polar_i)_{i=1}^n\}$, the goal of the pre-training task is to construct the knowledge-aware representation vectors $H = (h_1, \dots, h_n)^\top$ which can promote the downstream tasks in sentiment analysis. We devise a new supervised pre-training task called label-aware masked language model (LA-MLM), which introduces the sentence-level sentiment label l into the pre-training phase to capture the dependency between sentence-level language representation and individual words. It contains two separate sub-tasks: early fusion and late supervision.

3.3.1 Early Fusion

The purpose of early fusion is to recover the masked sequence conditioned on the sentence-level label, as shown in Figure 1. Assume that \hat{X}_k denotes the knowledge enhanced text sequence with some masked tokens, we can obtain the representation vectors with the input of \hat{X}_k and the sentence-level sentiment label l :

$$(h_{cls}^{EF}, h_1^{EF}, \dots, h_n^{EF}, h_{sep}^{EF}) = \text{Transformer}(\hat{X}_k, l)$$

where h_{cls}^{EF} and h_{sep}^{EF} are the hidden states of the special tokens [CLS] and [SEP]. The input embeddings of \hat{X}_k contains the embedding used in BERT (Devlin et al., 2019), the part-of-speech (POS) embedding and the word-level polarity embedding. Additionally, the embedding of the sentence-level sentiment label l is *early* added to the input embeddings. The model is required to predict the word, part-of-speech tag, and word-level polarity at the masked positions individually, thus the loss function is devised as follows:

$$\mathcal{L}_{EF} = - \sum_{t=1}^n m_t \cdot [\log P(x_t | \hat{X}_k, l) + \log P(pos_t | \hat{X}_k, l) + \log P(polar_t | \hat{X}_k, l)]$$

where m_t is an indicator function and equals to 1 iff x_t is masked. The prediction probabilities $P(x_t | \hat{X}_k, l)$, $P(pos_t | \hat{X}_k, l)$ and $P(polar_t | \hat{X}_k, l)$ are calculated based on the hidden state h_t^{EF} . This sub-task explicitly exerts the impact of the global sentiment label on the words and the linguistic knowledge of words, enhancing the ability of our model to explore the complex connection among them.

Task	Input Format	Output Hidden States
Sentence-level Sentiment Classification	[CLS] x_1, \dots, x_n [SEP]	$h_{[\text{CLS}]}$
Aspect Term Extraction	[CLS] $x_1 \dots x_n$ [SEP]	h_1, h_2, \dots, h_n
Aspect Term Sentiment Classification	[CLS] $a_1 \dots a_l$ [SEP] $x_1 \dots x_n$ [SEP]	$h_{[\text{CLS}]}$
Aspect Category Detection	[CLS] $x_1 \dots x_n$ [SEP]	$h_{[\text{CLS}]}$
Aspect Category Sentiment Classification	[CLS] $a_1 \dots a_l$ [SEP] $x_1 \dots x_n$ [SEP]	$h_{[\text{CLS}]}$
Text Matching	[CLS] $x_1 \dots x_n$ [SEP] $y_1 \dots y_m$ [SEP]	$h_{[\text{CLS}]}$

Table 1: Fine-tuning setting of SentiLARE on downstream tasks. Both $x_1 \dots x_n$ and $y_1 \dots y_m$ indicate the text sequences, while $a_1 \dots a_l$ denotes the aspect term / category sequence. The output hidden states are then used in the classification / regression layer.

3.3.2 Late Supervision

The late supervision sub-task aims to predict the sentence-level label and the word information based on the hidden states at [CLS] and masked positions respectively, as shown in Figure 1. Similar to early fusion, the representation vectors with the input of \hat{X}_k are obtained as follows:

$$(h_{cls}^{LS}, h_1^{LS}, \dots, h_n^{LS}, h_{sep}^{LS}) = \text{Transformer}(\hat{X}_k)$$

In this sub-task, the sentiment label l is used as the *late* supervision signal. Thus, the loss function to simultaneously predict the sentence-level sentiment label, words, and linguistic knowledge of words is shown as follows:

$$\begin{aligned} \mathcal{L}_{LS} = & -\log P(l|\hat{X}_k) - \sum_{t=1}^n m_t \cdot [\log P(x_t|\hat{X}_k) \\ & + \log P(pos_t|\hat{X}_k) + \log P(polar_t|\hat{X}_k)] \end{aligned}$$

where the sentence-level classification probability $P(l|\hat{X}_k)$ is calculated based on the hidden state h_{cls}^{LS} . This sub-task enables our model to capture the implicit relationship among the sentence-level representation at [CLS] and word-level linguistic knowledge at masked positions.

Since the two sub-tasks are separate, we empirically set the percentage of pre-training data provided for the late supervision sub-task as 20% and early fusion as 80%. As for the masking strategy, we increase the probability of masking words with positive / negative sentiment polarity from 15% in the setting of BERT to 30% because they are more likely to impact the sentiment of the whole text.

4 Experiment

4.1 Pre-training Dataset and Implementation

We adopted the Yelp Dataset Challenge 2019³ as our pre-training dataset. This dataset contains

³<https://www.yelp.com/dataset/challenge>

6,685,900 reviews with 5-class review-level sentiment labels. Each review consists of 127.8 words on average.

Since our method can adapt to all the BERT-style pre-trained models, we used RoBERTa (Liu et al., 2019) as the base framework to construct Transformer blocks in this paper, and discussed the generalization ability to other pre-trained models like BERT (Devlin et al., 2019) in the following experiment. The hyper-parameters of the Transformer blocks were set to be the same as RoBERTa-Base due to the limited computational power. Considering the high cost of training from scratch, we utilized the parameters of pre-trained RoBERTa⁴ to initialize our model. We also followed RoBERTa to use Byte-Pair Encoding vocabulary (Radford et al., 2019) whose size was 50,265. The maximum sequence length in the pre-training phase was 128, while the batch size was 400. We took Adam (Kingma and Ba, 2015) as the optimizer and set the learning rate to be 5e-5. The warmup ratio was 0.1. SentiLARE was pre-trained on Yelp Dataset Challenge 2019 for 1 epoch with label-aware masked language model, which took about 20 hours on 4 NVIDIA RTX 2080 Ti GPUs.

4.2 Fine-tuning Setting

We fine-tuned SentiLARE to the downstream tasks including sentence-level sentiment classification, aspect-level sentiment analysis, and general text matching tasks in our experiments. We adopted the fine-tuning settings in the existing work (Devlin et al., 2019; Xu et al., 2019), and showed the input format and the output hidden states of each task in Table 1. Note that the input embeddings in all the downstream tasks only contain BERT embedding, part-of-speech embedding and word-level polarity embedding. The hyper-parameters of fine-tuning on different datasets are reported in the Appendix.

⁴<https://github.com/pytorch/fairseq>

4.3 Baselines

We compared SentiLARE with general pre-trained models, task-specific pre-trained models and task-specific models without pre-training.

General Pre-trained Models: We adopted BERT (Devlin et al., 2019), XLNet (Yang et al., 2019), and RoBERTa (Liu et al., 2019) as general pre-trained baselines. These models achieve state-of-the-art performance on various NLP tasks.

Task-specific Pre-trained Models: We used BERT-PT (Xu et al., 2019), TransBERT (Li et al., 2019), and SentiBERT (Yin et al., 2020) as task-specific pre-trained baselines. Since TransBERT is not originally designed to deal with sentiment analysis tasks, we chose review-level sentiment classification on Yelp Dataset Challenge 2019 as the transfer task, and the downstream tasks in sentiment analysis as the target tasks.

Task-specific Models without Pre-training: We also chose some task-specific baselines without pre-training for corresponding tasks, including SC-SNN (Chen et al., 2019), DRNN (Wang, 2018), ML (Sachan et al., 2019) for sentence-level sentiment classification, DE-CNN (Xu et al., 2018) for aspect term extraction, CDT (Sun et al., 2019a) for aspect term sentiment classification, TAN (Movahedi et al., 2019) for aspect category detection, and AS-Capsules (Wang et al., 2019b) for aspect category sentiment classification.

We evaluated all the pre-trained baselines based on the codes and the model parameters provided by the original papers. For a fair comparison, all the pre-trained models were set to the base version, which possess a similar number of parameters (about 110M). The experimental results were presented with mean values over 5 runs. As for the task-specific baselines without pre-training, we re-printed the results on the corresponding benchmarks from the references.

4.4 Sentence-level Sentiment Classification

Dataset	Amount (Train/Valid/Test)	Length	# classes
SST	8,544 / 1,101 / 2,210	19.2	5
MR	8,534 / 1,078 / 1,050	21.7	2
IMDB	22,500 / 2,500 / 25,000	279.2	2
Yelp-2	504,000 / 56,000 / 38,000	155.3	2
Yelp-5	594,000 / 56,000 / 50,000	156.6	5

Table 2: Statistics of sentence-level sentiment classification (SSC) datasets.

We first evaluated our model on sentence-

level sentiment classification benchmarks including Stanford Sentiment Treebank (SST) (Socher et al., 2013), Movie Review (MR) (Pang and Lee, 2005), IMDB (Maas et al., 2011), and Yelp-2/5 (Zhang et al., 2015), which are widely used datasets at different scales. The detailed statistics of these datasets are shown in Table 2, which contain the amount of training / validation / test sets, the average length and the number of classes. Since MR, IMDB, and Yelp-2/5 don't have validation sets, we randomly sampled subsets from the training sets as the validation sets, and evaluated all the pre-trained models with the same data split.

Model	SST	MR	IMDB	Yelp-2	Yelp-5
SOTA-NPT	55.20 [#]	82.50 [#]	93.57 [†]	97.27 [‡]	69.15 [‡]
BERT	53.37	87.52	93.87	97.74	70.16
XLNet	56.33	89.45	95.27	97.41	70.23
RoBERTa	54.89	89.41	94.68	97.98	70.12
BERT-PT	53.24	87.30	93.99	97.77	69.90
TransBERT	55.56	88.69	94.79	96.73	69.53
SentiBERT	56.87	88.59	94.04	97.66	69.94
SentiLARE	58.59**	90.82**	95.71**	98.22**	71.57**

Table 3: Accuracy on sentence-level sentiment classification (SSC) benchmarks (%). SOTA-NPT means the state-of-the-art performance from the baselines without pre-training, where the results marked with [#], [†] and [‡] are re-printed from Chen et al. (2019), Sachan et al. (2019) and Wang (2018), respectively. ** indicates that our model significantly outperforms the best pre-trained baselines on the corresponding dataset (t-test, p -value < 0.01).

The results are shown in Table 3. We can observe that SentiLARE performs better than the baselines on the sentence-level sentiment classification datasets, thereby indicating the effectiveness of our knowledge-aware representation in sentiment understanding. Compared with vanilla RoBERTa, SentiLARE enhances the performance on all the datasets significantly. This demonstrates that for sentiment analysis tasks, linguistic knowledge can be used to enhance the state-of-the-art pre-trained model via the well-designed pre-training task.

4.5 Aspect-level Sentiment Analysis

Aspect-level sentiment analysis includes aspect term extraction, aspect term sentiment classification, aspect category detection, and aspect category sentiment classification. For aspect term based tasks, we chose SemEval2014 Task 4 for laptop (Lap14) and restaurant (Res14) domains (Pontiki et al., 2014) as the benchmarks, while for aspect category based tasks, we used SemEval2014 Task

Dataset	# sentences (Train/Test)	# terms (Train/Test)	# categories (Train/Test)	# sentiment classes	Amount of ATE (Train/Valid/Test)	Amount of ATSC (Train/Valid/Test)	Amount of ACD (Train/Valid/Test)	Amount of ACSC (Train/Valid/Test)
Lap14	3,045 / 800	2,358 / 654	-	3	1,338 / 150 / 422	2,163 / 150 / 638	-	-
Res14	3,041 / 800	3,693 / 1,134	3,711 / 1,025	3	1,871 / 150 / 606	3,452 / 150 / 1,120	2,891 / 150 / 800	3,366 / 150 / 973
Res16	2,000 / 676	-	2,507 / 859	3	-	-	1,850 / 150 / 676	2,150 / 150 / 751

Table 4: Statistics of aspect-level sentiment analysis datasets, where ATE / ATSC / ACD / ACSC indicates aspect term extraction / aspect term sentiment classification / aspect category detection / aspect category sentiment classification, respectively.

Task	ATE		ATSC				ACD		ACSC			
Dataset	Lap14	Res14	Lap14		Res14		Res14	Res16	Res14		Res16	
Model	F1	F1	Acc.	MF1.	Acc.	MF1.	F1	F1	Acc.	MF1.	Acc.	MF1.
SOTA-NPT	81.59 [‡]	-	77.19 [†]	72.99 [†]	82.30 [†]	74.02 [†]	90.61 [‡]	78.38 [‡]	85.00 ^b	73.53 ^b	-	-
BERT	83.22	87.68	78.18	73.11	83.77	76.06	90.48	72.59	88.35	80.40	86.55	71.19
XLNet	86.02	89.41	80.00	75.88	84.93	76.70	91.35	73.00	91.63	84.79	87.46	73.06
RoBERTa	87.25	89.55	81.03	77.16	86.07	79.21	91.69	77.89	90.67	83.81	88.38	76.04
BERT-PT	85.99	89.40	78.46	73.82	85.86	77.99	91.89	75.42	91.57	85.08	90.20	77.09
TransBERT	83.62	87.88	80.06	75.43	86.38	78.95	91.50	76.27	91.43	85.03	90.41	78.56
SentiBERT	82.63	88.67	76.87	71.74	83.71	75.42	91.67	73.13	89.68	82.90	87.08	72.10
SentiLARE	88.22*	91.15**	82.16*	78.70*	88.32**	81.63**	92.22	80.71**	92.97**	87.30**	91.29	80.00

Table 5: F1, accuracy (Acc.) and Macro-F1 (MF1.) on four aspect-level sentiment analysis tasks including aspect term extraction (ATE), aspect term sentiment classification (ATSC), aspect category detection (ACD) and aspect category sentiment classification (ACSC) (%). SOTA-NPT means the state-of-the-art performance from the baselines without pre-training, where the results marked with [‡], [†], [‡] and ^b are re-printed from Xu et al. (2018), Sun et al. (2019a), Movahedi et al. (2019) and Wang et al. (2019b), respectively. - means that the results are not reported in the references. * indicates that our model significantly outperforms the best pre-trained baselines on the corresponding dataset (t-test, p -value < 0.05), while ** means p -value < 0.01.

4 for restaurant domain (Res14) and SemEval2016 Task 5 for restaurant domain (Res16) (Pontiki et al., 2016). The statistics of these datasets are reported in Table 4. We followed the existing work (Xu et al., 2019) to leave 150 examples from the training sets for validation. Since the number of the examples with the *conflict* sentiment label is rather small, we adopted the same setting as the existing work (Tang et al., 2016; Xu et al., 2019) and dropped these examples in the aspect term / category sentiment classification task.

We present the results of aspect-level sentiment analysis in Table 5. We can see that SentiLARE outperforms the baselines on all the four tasks, and most of the improvement margins are significant. Interestingly, in addition to aspect-level sentiment classification, our model also performs well in aspect term extraction and aspect category detection. Since the aspect words are mostly nouns, part-of-speech tags may provide additional knowledge for the extraction task. In addition, the aspect terms can be signaled by neighboring sentiment words. This may explain why our knowledge-aware representation can help to extract (detect) the aspect term (category).

4.6 Ablation Test

To study the effect of the linguistic knowledge and the label-aware masked language model, we conducted ablation tests and presented the results in Table 6. Since the two sub-tasks are separate, the setting of -EF / -LS indicates that the pre-training data were all fed into the late supervision / early fusion sub-task, and -EF-LS denotes that the pre-training task is changed from label-aware masked language model to vanilla masked language model, while the input embeddings still include part-of-speech and word-level polarity embeddings. The -POS / -POL setting means that we removed the part of speech / word-level sentiment polarity in the input embeddings respectively, as well as in the supervision signals of the two sub-tasks. Obviously, -POS-POL indicates the complete removal of linguistic knowledge.

Results in Table 6 show that both the pre-training task and the linguistic knowledge contribute to the improvement over RoBERTa. Compared with early fusion, the late supervision sub-task plays a more important role in the classification tasks which depend on the global representation of the input sequence, such as SSC, ATSC, ACD and ACSC. Intuitively, the late supervision sub-task may learn a meaningful global representation at [CLS] by

Task	SSC	ATE	ATSC	ACD	ACSC
Dataset	SST	Res14	Res14	Res16	Res14
Model	Acc.	F1	Acc. MF1.	F1	Acc. MF1.
RoBERTa	54.89	89.55	86.07 79.21	77.89	90.67 83.81
SentiLARE	58.59	91.15	88.32 81.63	80.71	92.97 87.30
- EF	58.44	90.82	87.70 81.11	80.42	92.70 86.42
- LS	57.33	90.88	87.21 80.46	79.74	92.44 86.14
- EF - LS	56.91	90.74	86.95 79.71	78.92	91.32 84.73
- POS	58.15	90.94	87.98 81.38	80.27	92.51 86.61
- POL	57.95	90.63	87.64 81.34	79.40	92.46 86.30
- POS - POL	57.31	90.35	87.59 81.20	79.21	92.21 85.68

Table 6: Ablation test on sentiment analysis tasks. EF / LS / POS / POL denotes early fusion / late supervision / part-of-speech tag / word-level polarity, respectively.

simultaneously predicting the sentence-level sentiment labels and the word knowledge. Thus, it contributes more to the performance on these classification tasks.

As for the impact of linguistic knowledge, the performance of SentiLARE degrades more in the setting of removing the word-level sentiment polarity. This implies that the word-level polarity can help the pre-trained model more to derive the global sentiment in the classification tasks and signal neighboring aspects in the extraction task.

4.7 Analysis on Knowledge Acquisition

To investigate whether our proposed context-aware knowledge acquisition method can help construct knowledge-aware language representation, we compared the context-aware sentiment attention described in §3.2 with a context-free prior polarity acquisition algorithm (Guerini et al., 2013). This algorithm simply acquires a fixed sentiment polarity for each word with its part-of-speech tag by weighting the sentiment score of each sense with the reciprocal of the sense number, regardless of the variety of context. All the other parts of SentiLARE remain unchanged for comparison between these two knowledge acquisition methods.

Task	SSC	ATE	ATSC	ACD	ACSC
Dataset	SST	Res14	Res14	Res16	Res14
Model	Acc.	F1	Acc. MF1.	F1	Acc. MF1.
SentiLARE (w/ CASA)	58.59	91.15	88.32 81.63	80.71	92.97 87.30
SentiLARE (w/ CFPP)	56.86	90.54	88.05 81.26	79.00	92.72 86.59

Table 7: F1, accuracy (Acc.) and Macro-F1 (MF1.) of SentiLARE with context-aware sentiment attention (CASA) or context-free prior polarity (CFPP) on sentiment analysis tasks (%).

Results in Table 7 show that our context-aware

method performs better on all the tasks. This demonstrates that our context-aware attention mechanism can help SentiLARE model the sentiment roles of words across different contexts, thereby leading to better knowledge enhanced language representation.

4.8 Analysis on Knowledge Integration

To further demonstrate the importance of label-aware masked language model which deeply integrates linguistic knowledge into pre-trained models, we divided the test set of SST into two subsets according to the number of sentiment words (including positive and negative words determined by linguistic knowledge) in the sentences. Since there are 6.48 sentiment words on average in the sentences of SST’s test set, we partitioned the test set into two subsets: *SST-Less* contains the sentences with no more than 7 sentiment words, and *SST-More* includes the other sentences. Intuitively and presumably, the sentences with more sentiment words may include more complex sentiment expressions. We compared three models: RoBERTa which does not use linguistic knowledge, SentiLARE-EF-LS which simply augments input embeddings with linguistic knowledge as described in §4.6, and SentiLARE which deeply integrates linguistic knowledge via the pre-training task.

Model	SST-Less	SST-More
RoBERTa	55.61	53.52
SentiLARE-EF-LS	58.20 (+2.59)	54.49 (+0.97)
SentiLARE	59.72 (+4.11)	56.45 (+2.93)

Table 8: Accuracy of RoBERTa, SentiLARE-EF-LS, and SentiLARE on two subsets of SST’s test set (%).

Results in Table 8 show that SentiLARE-EF-LS already outperforms RoBERTa remarkably on SST-Less by simply augmenting input features with linguistic knowledge. However, on SST-More, SentiLARE-EF-LS only obtains marginal improvement over RoBERTa. As for our model, SentiLARE can consistently outperform RoBERTa and SentiLARE-EF-LS, and the margin between SentiLARE and SentiLARE-EF-LS is more evident on SST-More. This indicates that our pre-training task can help to integrate the local sentiment information reflected by word-level linguistic knowledge into global language representation, and facilitate the understanding of complex sentiment expressions.

4.9 Analysis on Generalization Ability

Generalization to Other Pre-trained Models:

To study whether the introduced linguistic knowledge and the proposed pre-training task can improve the performance of other pre-trained models besides RoBERTa, we chose BERT (Devlin et al., 2019) and ALBERT (Lan et al., 2020) as the base framework for evaluation. The experimental settings were set the same as SentiLARE based on RoBERTa.

Task	SSC	ATE	ATSC	ACD	ACSC
Dataset	SST	Res14	Res14	Res16	Res14
Model	Acc.	F1	Acc.	MF1.	F1
BERT	53.37	87.68	83.77	76.06	72.59
SentiLARE (w/BERT)	55.64	89.93	86.47	79.62	77.44
ALBERT	53.67	89.44	84.25	75.96	73.61
SentiLARE (w/ALBERT)	54.75	90.63	86.56	79.17	77.74

Table 9: F1, accuracy (Acc.) and Macro-F1 (MF1.) of BERT, ALBERT, and SentiLARE based on BERT / ALBERT on sentiment analysis tasks (%).

Results in Table 9 show that SentiLARE based on BERT / ALBERT outperforms vanilla BERT / ALBERT on the datasets of all the tasks, which demonstrates that our proposed method can adapt to different BERT-style pre-trained models to benefit the tasks in sentiment analysis.

Dataset	Amount (Train/Valid/Test)	Task Type	# classes
SCT	1,771 / 100 / 1,871	Classification	2
SICK	4,500 / 500 / 4,927	Classification	3
STSb	5,749 / 1,500 / 1,379	Regression	-

Table 10: Statistics of text matching datasets.

Generalization to Other NLP Tasks: Since sentiment is a common feature to improve the tasks of text matching (Cai et al., 2017; Li et al., 2019), we chose three text matching tasks to explore whether our sentiment-aware representations can also benefit these tasks, including story ending prediction, textual entailment, and semantic textual similarity. We evaluated RoBERTa and SentiLARE on the datasets of SCT (Mostafazadeh et al., 2016), SICK (Marelli et al., 2014), and STSb (Cer et al., 2017) for the three tasks respectively. The statistics of these datasets are reported in Table 10. We followed the existing work (Li et al., 2019) to pre-process the SCT dataset, and directly adopted the official version of SICK and STSb datasets.

Results in Table 11 show that SentiLARE can

Task	StoryEndPred	TextEntail	SemTextSim	
Dataset	SCT	SICK	STSb	
Model	Acc.	Acc.	Pear.	Spear.
RoBERTa	90.95	88.22	88.37	87.61
SentiLARE	92.47	89.42	89.15	88.50

Table 11: Accuracy (Acc.), Pearson (Pear.) and Spearman (Spear.) correlations of RoBERTa and SentiLARE on the tasks of story ending prediction, textual entailment, and semantic textual similarity (%).

enhance the performance of RoBERTa on these text matching tasks. This indicates that sentiment-related linguistic knowledge can be successfully integrated into the pre-trained language representation model to not only benefit the sentiment analysis tasks but also generalize to other sentiment-related NLP tasks. We will explore the generalization of our model to more NLP tasks as future work.

5 Conclusion

We present a novel pre-trained model called SentiLARE for sentiment analysis, which introduces linguistic knowledge from SentiWordNet via context-aware sentiment attention, and adopts label-aware masked language model to deeply integrate knowledge into BERT-style models through pre-training tasks. Experiments show that SentiLARE outperforms state-of-the-art language representation models on various sentiment analysis tasks, and thus facilitates sentiment understanding.

Acknowledgments

This work was jointly supported by the NSFC projects (key project with No. 61936010 and regular project with No. 61876096), and the Guoqiang Institute of Tsinghua University with Grant No. 2019GQG1. This work was also supported by Beijing Academy of Artificial Intelligence, BAAI. We thank THUNUS NEXt Joint-Lab for the support.

References

- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the International Conference on Language Resources and Evaluation*.
- Pierpaolo Basile, Annalina Caputo, and Giovanni Semeraro. 2014. An enhanced lesk word sense disambiguation algorithm through a distributional semantic model. In *25th International Conference on Computational Linguistics*, pages 1591–1600.

- Zheng Cai, Lifu Tu, and Kevin Gimpel. 2017. Pay attention to the ending: Strong neural baselines for the ROC story cloze task. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 616–622.
- Daniel M. Cer, Mona T. Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation*, pages 1–14.
- Shiyun Chen, Xin Lin, Yanghua Xiao, and Liang He. 2019. Sentiment commonsense induced sequential neural networks for sentiment classification. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 1021–1030.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Andrea Esuli and Fabrizio Sebastiani. 2006. SENTIWORDNET: A publicly available lexical resource for opinion mining. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation*, pages 417–422.
- Marco Guerini, Lorenzo Gatti, and Marco Turchi. 2013. Sentiment analysis: How to derive prior polarities from sentiwordnet. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1259–1269.
- Suchin Gururangan, Ana Marasovic, 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 the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 168–177.
- Minlie Huang, Qiao Qian, and Xiaoyan Zhu. 2017. Encoding syntactic knowledge in neural networks for sentiment classification. *ACM Transactions on Information Systems*, 35(3):26:1–26:27.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of International Conference on Learning Representations*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A lite BERT for self-supervised learning of language representations. In *8th International Conference on Learning Representations*.
- Zhongyang Li, Xiao Ding, and Ting Liu. 2019. Story ending prediction by transferable BERT. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pages 1800–1806.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv: 1907.11692*.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A SICK cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation*, pages 216–223.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James F. Allen. 2016. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849.
- Sajad Movahedi, Erfan Ghadery, Hesham Faili, and Azadeh Shakery. 2019. Aspect category detection via topic-attention network. *arXiv preprint arXiv:1901.01183*.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *43rd Annual Meeting of the Association for Computational Linguistics*, pages 115–124.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2227–2237.

- Matthew E Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A Smith. 2019. Knowledge enhanced contextual word representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 43–54.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad Al-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Véronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia V. Loukachevitch, Evgeniy V. Kotelnikov, Núria Bel, Salud María Jiménez Zafra, and Gülsen Eryigit. 2016. Semeval-2016 task 5: Aspect based sentiment analysis. In *Proceedings of the 10th International Workshop on Semantic Evaluation*, pages 19–30.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. Semeval-2014 task 4: Aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation*, pages 27–35.
- Qiao Qian, Minlie Huang, Jinhao Lei, and Xiaoyan Zhu. 2017. Linguistically regularized LSTM for sentiment classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1689.
- Qiao Qian, Bo Tian, Minlie Huang, Yang Liu, Xuan Zhu, and Xiaoyan Zhu. 2015. Learning tag embeddings and tag-specific composition functions in recursive neural network. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing*, pages 1365–1374.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. In *OpenAI Technical Report*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. In *OpenAI Technical Report*.
- Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 3980–3990.
- Devendra Singh Sachan, Manzil Zaheer, and Ruslan Salakhutdinov. 2019. Revisiting LSTM networks for semi-supervised text classification via mixed objective function. In *The Thirty-Third AAAI Conference on Artificial Intelligence*, pages 6940–6948.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Kai Sun, Richong Zhang, Samuel Mensah, Yongyi Mao, and Xudong Liu. 2019a. Aspect-level sentiment analysis via convolution over dependency tree. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 5678–5687.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019b. ERNIE: enhanced representation through knowledge integration. *arXiv preprint arXiv: 1904.09223*.
- Duyu Tang, Bing Qin, and Ting Liu. 2016. Aspect level sentiment classification with deep memory network. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 214–224.
- Zhiyang Teng, Duy-Tin Vo, and Yue Zhang. 2016. Context-sensitive lexicon features for neural sentiment analysis. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1629–1638.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations*.
- Baoxin Wang. 2018. Disconnected recurrent neural networks for text categorization. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 2311–2320.
- Yequan Wang, Aixin Sun, Minlie Huang, and Xiaoyan Zhu. 2019b. Aspect-level sentiment analysis using as-capsules. In *The World Wide Web Conference*, pages 2033–2044.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 347–354.
- Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2018. Double embeddings and cnn-based sequence labeling for aspect extraction. In *Proceedings of the 56th Annual*

Meeting of the Association for Computational Linguistics, pages 592–598.

Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2019. BERT post-training for review reading comprehension and aspect-based sentiment analysis. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2324–2335.

Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2020. Dombert: Domain-oriented language model for aspect-based sentiment analysis. *arXiv preprint arXiv: 2004.13816*.

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.

Da Yin, Tao Meng, and Kai-Wei Chang. 2020. Sentibert: A transferable transformer-based architecture for compositional sentiment semantics. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3695–3706.

Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*, pages 649–657.

Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. ERNIE: enhanced language representation with informative entities. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 1441–1451.

A Experimental Details

A.1 Hyper-parameter Setting

We provided the hyper-parameter search space during pre-training in Table 12. Grid search was used to select hyper-parameters, and the selection criterion was the classification accuracy on the validation set when we fine-tuned the pre-trained model on SST.

We also provided the detailed setting of hyper-parameters during fine-tuning on the datasets of sentiment analysis, including hyper-parameter search space in Table 13 and best assignments in Table 14. Note that we used HuggingFace’s Transformers⁵ to implement our model, so all the hyper-parameters we reported were consistent with the codes of HuggingFace’s Transformers. We utilized manual search to select the best hyper-parameters

⁵<https://github.com/huggingface/transformers>

Hyper-parameter	Search Space
Percentage of Data in EF	<i>choice</i> [0,0.2,0.4,0.6,0.8,1]
Masking Probability (Sentiment Words)	<i>choice</i> [0.2,0.3,0.4,0.5]
Batch Size	<i>choice</i> [400,512]
Training Epoch	<i>choice</i> [1,2]
Learning Rate	5e-5
Warmup Ratio	0.1
Sequence Length	128
Maximum Gradient Norm	1.0
Optimizer	Adam
Epsilon (for Adam)	1e-8

Table 12: Hyper-parameter search space of SentiLARE during pre-training. *choice* indicates that the listed numbers will be chosen with the same probability.

during fine-tuning. The number of hyper-parameter search trials for each dataset was 20. We used accuracy as our criterion for selection on all the sentiment analysis tasks except aspect term extraction and aspect category detection. For these two tasks, F1 was adopted as the selection criterion.

Hyper-parameter	Search Space
Learning Rate	<i>choice</i> [1e-5,2e-5,3e-5,4e-5,5e-5]
Training Epoch	<i>uniform-integer</i> [3,8]
Warmup Step	<i>uniform-integer</i> [0,total_step*0.2]
Batch Size	<i>choice</i> [12,16,24,32]
Sequence Length	<i>choice</i> [128,256,512]
Maximum Gradient Norm	1.0
Optimizer	Adam
Epsilon (for Adam)	1e-8

Table 13: Hyper-parameter search space of SentiLARE on the downstream sentiment analysis tasks. *uniform-integer* means the integers in the interval can be selected uniformly. In the search space of warmup step, total_step denotes the total training steps on different datasets.

A.2 Results on Validation Sets

In addition to the performance on the test set of each dataset which has been reported in the main paper, we also provided the validation performance on the datasets of sentence-level and aspect-level sentiment analysis in Table 15. As mentioned in Appendix A.1, accuracy and F1 were used to select the best hyper-parameters, so we reported the validation performance of all the pre-trained models on these metrics.

A.3 Runtime

The runtime of fine-tuning on different datasets of sentiment analysis was reported in Table 16. We tested all the pre-trained models on 4 NVIDIA RTX 2080 Ti GPUs.

Task	SSC					ATE		ATSC		ACD		ACSC	
Dataset	SST	MR	IMDB	Yelp-2	Yelp-5	Lap14	Res14	Lap14	Res14	Res14	Res16	Res14	Res16
Learning Rate	2e-5	3e-5	2e-5	2e-5	2e-5	3e-5	3e-5	3e-5	3e-5	3e-5	3e-5	3e-5	3e-5
Training Epoch	3	4	3	3	3	4	4	8	8	4	6	8	8
Warmup Step	100	20	100	12,600	8,500	0	0	0	150	100	0	0	60
Batch Size	12	24	24	12	12	12	12	16	16	16	16	16	32
Sequence Length	256	256	512	512	512	128	128	128	128	128	128	128	128

Table 14: Best assignments of hyper-parameters on the sentiment analysis tasks.

Task	SSC					ATE		ATSC		ACD		ACSC	
Dataset	SST	MR	IMDB	Yelp-2	Yelp-5	Lap14	Res14	Lap14	Res14	Res14	Res16	Res14	Res16
Model	Acc.	Acc.	Acc.	Acc.	Acc.	F1	F1	Acc.	Acc.	F1	F1	Acc.	Acc.
BERT	52.77	87.05	93.72	97.77	70.56	81.62	83.48	80.67	77.33	86.96	72.12	84.67	86.00
XLNet	54.13	89.05	95.20	97.48	70.69	85.13	87.80	82.00	82.00	86.63	71.72	84.00	82.67
RoBERTa	54.13	89.71	94.68	98.06	70.65	88.75	88.09	84.67	82.67	89.21	77.62	86.00	85.33
BERT-PT	52.32	88.00	93.88	97.82	70.34	86.55	86.39	84.67	80.00	87.65	79.02	84.67	88.00
TransBERT	53.41	88.57	94.72	96.82	70.15	84.52	85.20	80.67	79.33	87.39	77.14	87.33	88.67
SentiBERT	54.50	88.76	94.12	97.83	70.44	83.23	83.48	82.67	81.33	86.90	73.65	84.67	84.00
SentiLARE	55.04	90.07	95.96	98.26	72.14	86.07	88.61	83.33	84.67	88.37	80.54	88.67	89.33

Table 15: Accuracy (Acc.) and F1 on the validation sets of sentiment analysis benchmarks (%).

Task	SSC					ATE		ATSC		ACD		ACSC	
Dataset	SST	MR	IMDB	Yelp-2	Yelp-5	Lap14	Res14	Lap14	Res14	Res14	Res16	Res14	Res16
BERT	22	16	42	1,389	1,492	5	5	15	23	9	9	10	9
XLNet	28	23	158	3,412	3,607	5	8	17	25	10	11	12	8
RoBERTa	23	18	65	1,674	1,930	5	6	16	26	9	8	10	9
BERT-PT	25	18	53	1,426	1,611	4	6	17	26	8	7	10	6
TransBERT	21	19	42	1,413	1,586	5	5	12	17	10	9	10	7
SentiBERT	16	16	53	1,336	1,503	3	4	9	13	7	5	8	5
SentiLARE	26	20	56	1,539	1,706	7	8	13	20	12	13	10	7

Table 16: Average runtime of fine-tuning on different datasets for each pre-trained model (minutes).

B Case Study

To intuitively show that SentiLARE can integrate linguistic knowledge into pre-trained models to promote sentiment analysis, we provided a case, and visualized the classification probability of all the prefix subsequences truncated at each position in Figure 2. For example, the prediction probability at the position of *small* was acquired by the hidden state of [CLS] obtained by RoBERTa / SentiLARE with the input “[CLS] The restaurant is small [SEP]”. Both RoBERTa and SentiLARE were fine-tuned on the Yelp-2 dataset.




Figure 2: Visualization of the classification probability of RoBERTa and SentiLARE fine-tuned on Yelp-2. The bar indicates the output distribution of RoBERTa and SentiLARE with the input of the prefix subsequence at each position.

Compared with RoBERTa, our model enhanced with word-level linguistic knowledge can successfully capture the sentiment shift caused by the word *change* in this sentence, thereby determining the correct sentence-level sentiment label.

C Analysis on Textual Similarity in Knowledge Acquisition

Task	SSC	ATE	ATSC	ACD	ACSC
Dataset	SST	Res14	Res14	Res16	Res14
Model	Acc.	F1	Acc. MF1.	F1	Acc. MF1.
SentiLARE (w/SBERT)	58.59	91.15	88.32 81.63	80.71	92.97 87.30
SentiLARE (w/WordVec)	57.41	90.66	88.10 81.27	80.07	92.75 86.87

Table 17: F1, accuracy (Acc.) and Macro-F1 (MF1.) of SentiLARE with Sentence-BERT (SBERT) or word vectors (WordVec) on sentiment analysis tasks (%).

Since Sentence-BERT is costly to calculate the textual similarity between contexts and glosses, we compared it with another lighter textual similarity algorithm (Basile et al., 2014) which computes the representation vectors of sentences by averaging

the embedding vectors of their constituent words. We used 300-dimensional GloVe⁶ as the word vectors to obtain textual similarity.

Results in Table 17 show that Sentence-BERT performs better on all the tasks. Nevertheless, static word vectors are more computationally efficient in linguistic knowledge acquisition with acceptable performance drop.

⁶<https://nlp.stanford.edu/projects/glove/>