

SoftMCL: Soft Momentum Contrastive Learning for Fine-grained Sentiment-aware Pre-training

Jin Wang[†], Liang-Chih Yu^{‡,*} and Xuejie Zhang[†]

[†]School of Information Science and Engineering, Yunnan University, Kunming, China

[‡]Department of Information Management, Yuan Ze University, Taoyuan, China
{wangjin, xjzhang}@ynu.edu.cn, lcyu@saturn.yzu.edu.tw

Abstract

The pre-training for language models captures general language understanding but fails to distinguish the affective impact of a particular context to a specific word. Recent works have sought to introduce contrastive learning (CL) for sentiment-aware pre-training in acquiring affective information. Nevertheless, these methods present two significant limitations. First, the compatibility of the GPU memory often limits the number of negative samples, hindering the opportunities to learn good representations. In addition, using only a few sentiment polarities as hard labels, e.g., positive, neutral, and negative, to supervise CL will force all representations to converge to a few points, leading to the issue of latent space collapse. This study proposes a soft momentum contrastive learning (SoftMCL) for fine-grained sentiment-aware pre-training. Instead of hard labels, we introduce valence ratings as soft-label supervision for CL to fine-grained measure the sentiment similarities between samples. The proposed SoftMCL is conducted on both the word- and sentence-level to enhance the model’s ability to learn affective information. A momentum queue was introduced to expand the contrastive samples, allowing storing and involving more negatives to overcome the limitations of hardware platforms. Extensive experiments were conducted on four different sentiment-related tasks, which demonstrates the effectiveness of the proposed SoftMCL method. The code and data of the proposed SoftMCL is available at: <https://www.github.com/wangjin0818/SoftMCL/>.

Keywords: Sentiment-aware pre-training, contrastive learning, momentum queue, valence-arousal ratings

1. Introduction

Pre-trained language models (PLMs), such as BERT (Devlin et al., 2019), XLNet (Yang et al., 2019), RoBERTa (Liu et al., 2019), and DeBERTa (He et al., 2021), have proven their powerful ability to learn representations for natural language understanding and generation tasks. To learn language usage, these models were first trained on large-scale unlabeled corpora by unsupervised tasks, such as masked language model (MLM) and next sentence prediction (NSP). Then, they can be transferred to another task by supervised training on a smaller task-specific dataset. The great success of those PLMs is attributed to the practical pre-training tasks.

The core idea of pre-training is to generate similar contextual representations for words or sentences with similar contexts. Although the training paradigm works well for semantic-oriented applications with supervision for the role of language usage, they fail to distinguish the affective impact of a particular context to a specific word since certain words may have special semantics and sentiments in specific contexts (Agrawal et al., 2018). For example, the emotion of *The battery life is long* is positive, but *It takes a long time to focus* is negative, even though both sentences contain the word *long*.

Similarly, *The scope of his book is ambitious*, and *The government’s decisions to begin the ambitious plans which cost a lot*, contain the same *ambitious* vocabulary but clearly express the opposite attitude. As a result, PLMs are still sub-optimal in sentiment-related tasks due to the ignorance of the affective information in pre-training. As reported by Kassner and Schütze (2020), It is still hard for PLMs to handle words with contradiction sentiment or negation expression, which are critical in sentiment analysis tasks.

To assist PLMs in obtaining affective information, early exploration works have sought to use either *sentiment-aware pre-training* (Abdalla et al., 2019; Fu et al., 2018) or *sentiment refinement* (Utsumi, 2019; Yu et al., 2018a,b). Both methods aim to obtain representations that are similar in semantics and sentiment. The *sentiment-aware pre-training* accomplishes this goal by introducing external affective information as post-training supervision on word- or sentence-level, such as token sentiments and emoticons (Zhou et al., 2020), aspect word (Tian et al., 2020), linguistic knowledge (Ke et al., 2020), and implicit sentiment information (Li et al., 2021). The external sentiment knowledge can be introduced by performing word- or sentence-level polarity classification on the masked words (Ke et al., 2020). Nevertheless, learning word sentiment cannot help the model understand the sentiment intention of the whole sentence. Since the

*Corresponding author

expressed sentiment of a sentence is not simply the sum of the polarities or the intensity of its constituent words.

The *sentiment refinement* model adjusts the vector of each affective word according to a given sentiment lexicon (Utsumi, 2019; Yu et al., 2018a,b). It can be used to adjust the representations of words so that they can be closer to semantically and sentimentally similar words and further away from sentimentally dissimilar ones. Such a method is applied as a post-processing step for already trained representations. An obvious risk is that the improper incorporation of affective information may excessively disperse the distribution of the semantic space.

Existing studies suggest introducing contrastive learning (CL) and using the same idea as *sentiment refinement* to incorporate affective information. The idea involves training a model to pull together an anchor and a *positive* sample in the latent space and push apart the anchor from many *negative* samples. Semantic learning works (Gao et al., 2021) applied self-supervised CL, where a positive pair often consists of data augmentations of the sample, and negative pairs are formed by the anchor and randomly chosen samples in the batch. Unfortunately, this is unsuitable for sentiment-aware pre-training. If two words with the same sentiment polarity come from different samples in the same batch, they will be treated as negative sample pairs for CL, as the words *clumsy* and *bad* in Figure 1(a).

By leveraging sentiment information, recent studies (Fan et al., 2022; Li et al., 2021) applied supervised CL for sentiment-aware pre-training on both word- and sentence-level. The idea is transferred to that normalized representation from the words or sentences with the same sentiment polarities are pulled closer together than the ones with different sentiment polarities. By introducing external knowledge, such as sentiment-annotated lexicon or corpora, these models draw positives from samples with the same polarity as the anchor rather than data augmentation of the anchor, as shown in Figure 1(b).

A few practical issues should be considered when applying CL, especially for sentiment-aware pre-training. First is the batch size. As in most circumstances, larger batch size is better since that will provide more diverse and complicated negative samples, which is crucial for learning good representations (Chen et al., 2022). Nevertheless, the compatibility of the GPU memory often limits the batch size. Another important consideration is the quality of the negative samples. Hard negatives improve the representation quality learned by supervised CL if negative samples can be strictly classified according to the sentiment labels. It is evident that *good* and *excellent* are both positive, while the latter is stronger than the former. By defining only

a few sentiment categories, e.g., positive, neutral, and negative, supervised CL can also cause all representations to converge to a few points, leading to the issue of latent space collapse.

This study proposed a soft momentum contrastive learning (SoftMCL) for fine-grained sentiment-aware pre-training. In contrast to hard labels of sentiment polarities, we introduce valence ratings¹ from the extended version of affective norms of English words (E-ANEW) (Warriner et al., 2013) and EmoBank (Buechel and Hahn, 2016, 2017) as fine-grained soft-labels supervisions for CL. SoftMCL does not strictly distinguish between positive and negative samples, but measures the cross-entropy between the sentiment similarities in continuous affective space and distribution of semantic representations of these samples. Further, the proposed SoftMCL is conducted on both the word- and sentence-level to enhance the ability to learn affective information. A momentum queue was introduced to expand the contrastive samples, allowing storing and involving more negatives to overcome the limitations of hardware platforms.

Extensive experiments were conducted on four different sentiment-related tasks to evaluate the effectiveness of the proposed SoftMCL method. Empirical results show that SoftMCL performs better than other sentiment-aware pre-training methods.

The rest of the paper is organized as follows. Section 2 provides the preliminary knowledge of contrastive learning. Section 3 describes the proposed sentiment-aware momentum contrastive learning for fine-grained sentiment pre-training on both token- and sentence-level. Section 4 summarizes the comparative results against several baselines. Section 5 introduces the related works for affective state representation and sentiment pre-training. Conclusions are finally drawn in Section 6.

2. Preliminary

2.1. Valence-Arousal Ratings

Affective computing research offers two approaches to representing an emotional state: categorical and dimensional (Calvo and Mac Kim, 2013). The categorical approach models affective states as several distinct classes, such as binary (positive and negative), Ekman's six fundamental emotions (anger, happiness, fear, sadness, disgust, and surprise) (Ekman, 1992), and Plutchik's eight fundamental emotions (which include Ekman's six as well as trust and anticipation) (Plutchik, 1991). A set of emotion categories must be established before

¹The valence ratings in Figure 1 are clumsy (4.14), bad (3.24), good (7.89), early (5.26)

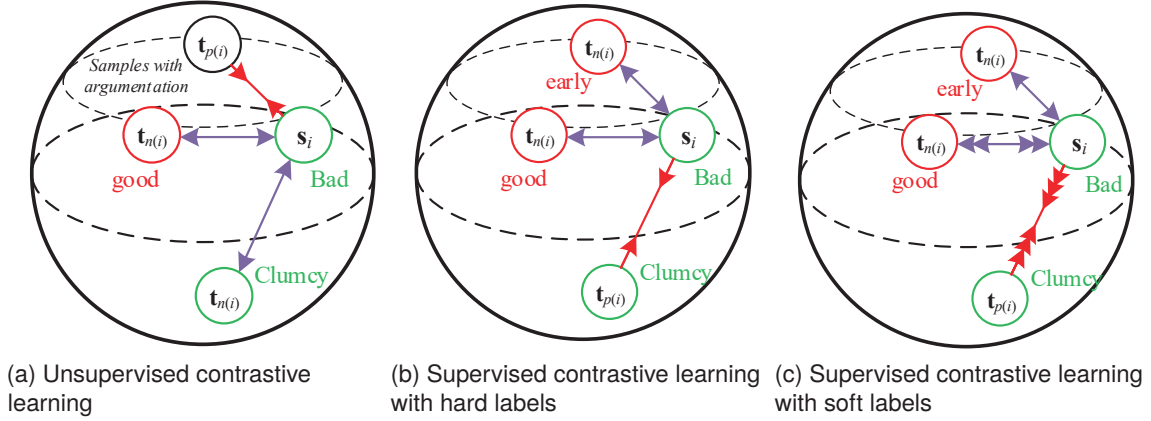


Figure 1: The conceptual diagram of using different contrastive learning strategies for sentiment-aware pre-training. (a) The self-supervised CL contrasts a single positive for each anchor (i.e., an augmentation of the anchor) against a set of negative consisting of the entire remainder of the batch. (b) The supervised CL contrasts the set of samples with same polarity as positives against the negatives from the remainder of the batch. (c) The proposed SoftMCL introducing external affective supervision to contrast the set of all samples according to the fine-grained distance of valence ratings between samples.

classification. Custom-defined categories, however, typically cannot cover the complete emotional spectrum. As a result, it's possible to classify some affective states into an undefined category. Researchers may describe the same category under different names due to cultural or other differences, making it challenging to share categories. Based on this representation, several methods have been researched to provide practical applications, such as customer review analysis, mental illnesses identification, hotspot detection and forecasting, question answering, social text analysis, and financial market prediction.

These are the three fundamental limitations for discrete categorical encoding of sentiment: 1) a set of emotion categories must be established before classification. Custom-defined categories, however, typically cannot cover the complete emotional spectrum. As a result, it's possible to classify some affective states into an undefined category. 2) Research teams may describe the same category under different names due to cultural or other differences, making it challenging to share categories. 3) The exceptional categories in particular domains must be re-defined when the field of application changes.

Continuous dimensional representations could be preferable because they express the affective state in a low-dimensional (2- or 3-dimensional) space. The most widely used method is based on Russell's valence-arousal space (Russell, 1980), which can precisely describe the affective state in a 2-dimensional continuous space (Lee et al., 2022). Valence and arousal are both real values. Valence is the pleasantness (both positive and negative) quality of a feeling. On the other side, arousal

describes the intensity of the emotion's activation (such as excitement or calmness). This approach can represent any affective state as a point in the valence-arousal coordinate plane, effectively avoiding inconsistency and incompleteness in category definitions in a discrete representation and providing more intelligent and fine-grained sentiment applications (Wang et al., 2019, 2020).

2.2. Contrastive Learning

Contrastive learning is a self-supervised learning paradigm that captures inherent patterns and context in data without manual annotations (Gao et al., 2021). The powerful representation learning ability of PLMs owes to the pre-training tasks in large-scale unlabeled data. Using contrastive learning, the PLMs can be trained by distinguishing the differences between samples. Given an anchor sample, a positive sample, and several negative samples, similarities are measured pairwise between the given anchor and the rest. Data augmentation usually produces a positive sample, meaning it belongs to the same class as the anchor sample. Conversely, negative samples indicate those not in the same class as the anchor. By measuring the similarity scores, the training objective of contrastive learning is,

$$\mathcal{L}_{CL} = - \sum_{i \in I} \log \frac{\exp(\mathbf{h}_i^\top \cdot \mathbf{h}_{p(i)} / \tau)}{\sum_{k \in A(i)} \exp(\mathbf{h}_i^\top \cdot \mathbf{h}_k / \tau)} \quad (1)$$

where \mathbf{h}_i denotes the hidden representation of the i th sample, I is the set of samples in a batch, \cdot is a dot product, τ is a temperature parameter used to

control the density, $p(i)$ represents its positive sample, and $A(i)$ is the set that contains the positive sample and the negative samples.

2.3. Supervised Contrastive Learning

In addition to self-supervised learning, the prior knowledge of the labeled data can be incorporated. Since the anchor label is known, the label annotation can identify the positive and negative samples (Fan et al., 2022; Li et al., 2021). Therefore, the training objective of contrastive learning can be generalized as follows,

$$\mathcal{L}_{\text{SCL}} = - \sum_{i \in I} \sum_{p \in P(i)} \frac{1}{|P(i)|} \log \frac{\exp(\mathbf{h}_i^\top \cdot \mathbf{h}_p / \tau)}{\sum_{k \in A(i)} \exp(\mathbf{h}_i^\top \cdot \mathbf{h}_k / \tau)} \quad (2)$$

where $P(i)$ indicates all positives corresponding to anchor i retrieved by the label annotation within a batch, and $|P(i)|$ is the number of positive samples. Using supervised learning, the PLMs are trained to pull the same class together while pushing apart clusters of samples from different classes.

In the objective of CL, the labels of positive (similar) pairs are set to 1, while the ones of negative (dissimilar) pairs are set to 0. Such hard labels lead to the issue of contrastive bias. Although two words have opposite emotional polarities, they may not differ significantly in valence.

3. Soft Momentum Contrastive Learning

Figure 2 shows the overall architecture of the proposed SoftMCL for fine-grained sentiment-aware pre-training. We introduce the affective ratings from the E-ANEW (Warriner et al., 2013) and EmoBank (Buechel and Hahn, 2016, 2017). The difference in valence ratings between samples thus measures the sentiment similarity. The SoftMCL is performed on both token- and sentence-level to pull together the sentimentally similar samples but push away the dissimilar ones. In addition to the sentiment-aware pre-training, we introduce a momentum queue to expand the contrastive samples. The pre-trained encoder could be fine-tuned for the downstream applications.

3.1. Sentiment Similarity

All emotions can be annotated in continuous VA space as a real value in $[1, 9]$ for valence, arousal, and dominance. Values 1, 5, and 9 denote valence’s most negative, neutral, and positive affective states.

Given an input sentence $\mathbf{X} = \{x_1, x_2, \dots, x_n\}$ where n denotes the sequence length, each $x_i \in \mathbf{X}$ denotes the i th word. By adding a special token

$[\text{CLS}]$, the input \mathbf{X} is transformed into representation vectors, denoted as,

$$\{\mathbf{h}_{[\text{CLS}]}, \mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n\} = \text{Enc}(\{\mathbf{x}_{[\text{CLS}]}, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}; \theta) \quad (3)$$

where $\text{Enc}(\cdot)$ is the Transformer encoder parameterized by θ . The associated affective ratings with both the sentence and words are obtained as, $Y = \{y_{[\text{CLS}]}, y_1, y_2, \dots, y_n\}$ where $y_{[\text{CLS}]}$ is the valence of the whole sentence, and the length n in Y should be equal to the length n in \mathbf{X} . If an affective word was split to several tokens, all tokens should share the same valence as the word. Additionally, the y_i is an annotated valence in a given affective lexicon. If a word x_i does not appear in the affective lexicon, its corresponding y_i will be assigned to 0, such that it will be masked during the word-level contrastive learning.

Based on the valence ratings, the sentiment similarity between two tokens or sentences can be defined as a normalized absolute similarity,

$$\Delta(\mathbf{h}_1, \mathbf{h}_2) = 1 - \frac{|y_1 - y_2|}{y_{\max} - y_{\min}} \quad (4)$$

where y_1 and y_2 are the valence ratings of the tokens x_1 and x_2 or the sentences \mathbf{X}_1 and \mathbf{X}_2 .

3.2. Sentiment-aware Contrastive Learning

The previous study applies unsupervised CL to contrast sentences within a batch. We perform both word- and sentence-level CL to enrich the sentiment information for the encoder. Using sentiment similarity of valence between samples, we do not strictly distinguish between positive and negative samples. Thus, the objective of sentiment-aware contrastive learning aims to measure the cross-entropy of the sentiment similarity distribution between samples and the similarity distribution of semantic representations, denoted as,

$$\mathcal{L}_{\text{SentiCL}} = - \sum_{i \in I} \sum_{j \in B(i)} \frac{\Delta(\mathbf{h}_i, \mathbf{h}_j)}{\sum_{l \in B(i)} \Delta(\mathbf{h}_i, \mathbf{h}_l)} \log \frac{\exp(\mathbf{h}_i^\top \cdot \mathbf{h}_j / \tau)}{\sum_{k \in B(i)} \exp(\mathbf{h}_i^\top \cdot \mathbf{h}_k / \tau)} \quad (5)$$

where $B(i)$ is the batch size. For token-level CL, we randomly sampled 256 tokens belonging to the affective words appearing in the E-ANEW lexicon. For sentence-level CL, we extract hidden representation corresponding to $[\text{CLS}]$ token, i.e., $\mathbf{h}_{[\text{CLS}]}$ to measure the cross-entropy of the sentiment and semantic similarity distribution of the in-batch sentences.

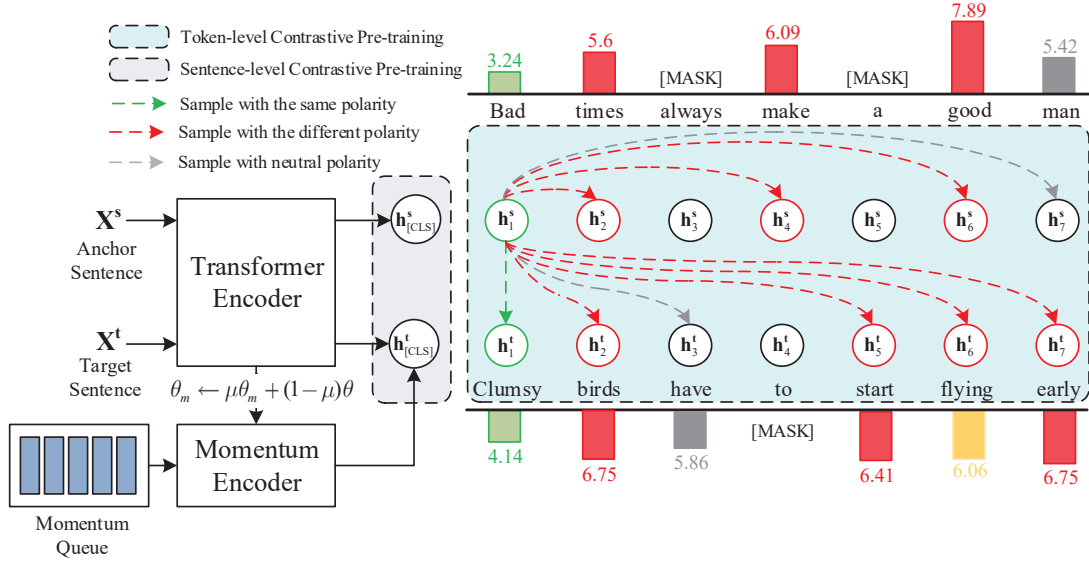


Figure 2: The overall architecture of the proposed SoftMCL.

3.3. Momentum Contrastive Learning

The batch size setting limits the number of contrastive samples. For sentence-level CL, each sample could only be paired with other $B(i) - 1$ samples to create additional training patterns. To further expand the contrastive samples in batch, we proposed a momentum queue strategy, which uses the same idea as the MoCo (He et al., 2020) objective. We applied a momentum encoder $\text{Enc}_m(\cdot)$ to learn the sequence of hidden representations of the recent samples and store them within the affective ratings in a momentum queue \mathcal{Q} . These stored samples can expand the number of contrastive samples. For implementation, the momentum encoder $\text{Enc}_m(\cdot)$ is copied from the initial encoder $\text{Enc}(\cdot)$. The parameters θ_m were updated from the parameter θ with a ratio μ , denoted as,

$$\theta_m \leftarrow \mu\theta_m + (1 - \mu)\theta \quad (6)$$

Based on the momentum queue \mathcal{Q} and encoder $\text{Enc}_m(\cdot)$, there will be $Q(i)$ samples that belong to \mathcal{Q} ; thus, we obtain $B(i) + Q(i)$ contrastive samples. The loss function still measures cross-entropy between the distribution of sentiment similarities and the distribution of semantic representations,

$$\mathcal{L}_{\text{MCL}} = - \sum_{i \in I} \left[\sum_{j \in B(i)} \frac{\Delta(\mathbf{h}_i, \mathbf{h}_j)}{\sum_{l \in \{B(i), Q(i)\}} \Delta(\mathbf{h}_i, \mathbf{h}_l)} \log \frac{\exp(\mathbf{h}_i^\top \cdot \mathbf{h}_j / \tau)}{\sum_{k \in \{B(i), Q(i)\}} \exp(\mathbf{h}_i^\top \cdot \mathbf{h}_k / \tau)} \right. \\ \left. + \sum_{m \in Q(i)} \frac{\Delta(\mathbf{h}_i, \mathbf{h}_m)}{\sum_{l \in \{B(i), Q(i)\}} \Delta(\mathbf{h}_i, \mathbf{h}_l)} \log \frac{\exp(\mathbf{h}_i^\top \cdot \mathbf{h}_m / \tau)}{\sum_{k \in \{B(i), Q(i)\}} \exp(\mathbf{h}_i^\top \cdot \mathbf{h}_k / \tau)} \right] \quad (7)$$

Dataset	Train	Dev	Test	Length	#C
SST-2	8,544	1,101	2,210	19.2	2
SST-5	6,920	872	1,821	19.2	5
MR	8,534	1,078	1,050	21.7	2
CR	2,025	675	675	20.1	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
Emobank	6,195	2,065	2,065	17.6	-
Facebook	1,736	579	579	19.8	-
Lap14	3,452	150	676	30.2	3
Rest14	2,163	150	638	25.6	3

Table 1: Statistics of datasets used in the experiments. #C indicates the number of the labels.

3.4. Training Objectives

In addition to sentiment-aware pre-training, masked language modeling (MLM) was also implemented for pre-training. Like the previous settings, 15% of words were first masked. Further, we masked affective words in the input sentence using the E-ANEW lexicon with a specific ratio. The MLM objective is to predict the masked words within the given contexts. Then, SoftMCL was performed to distinguish the valence difference between samples. The training objective is,

$$\mathcal{L} = \mathcal{L}_{\text{MLM}} + \lambda_1 \mathcal{L}_{\text{MCL}}^{\text{word}} + \lambda_2 \mathcal{L}_{\text{MCL}}^{\text{sent}} \quad (8)$$

where λ_1 and λ_2 balance the contribution of different subtasks.

For fine-tuning, the pre-trained model generates a sequence of hidden representations using Eq.(3). The aspect term was appended to the sequence with the [SEP] token for aspect-level sentiment classification. The hidden representation of the [CLS] token was used to predict the sentiment

Model	Word		Phrase	
	$k \uparrow$	$\rho \uparrow$	$k \uparrow$	$\rho \uparrow$
General pre-trained models				
BERT	0.638	0.841	0.624	0.829
XLNet	0.644	0.852	0.637	0.833
RoBERTa	0.656	0.868	0.642	0.849
DeBERTa	0.667	0.866	0.644	0.852
Sentiment-aware pre-trained models				
BERT-PT	0.704	0.880	0.688	0.877
SentiBERT	0.699	0.886	0.692	0.882
SentiLARE	0.712	0.894	0.702	0.891
Proposed models with ablation study				
SoftMCL	0.778	0.928	0.775	0.922
w/o WP	0.698	0.878	0.692	0.874
w/o SP	0.702	0.889	0.685	0.865
w/o MoCL	0.724	0.895	0.713	0.904

Table 2: Comparative results of different models on phrase-level intensity prediction task.

intensity or polarity of the input sample.

4. Experiments

4.1. Tasks and Datasets

Sentiment-aware Pre-training. The proposed SoftMCL was pre-trained using the same dataset as RoBERTa (Liu et al., 2019) and DeBERTa (He et al., 2021), which is a combination of Wikipedia, Bookcorpus, CCNews, Stories, and OpenWebText. We sampled 2 million sentences with a maximum length of 128 for word-level pre-training. We introduce the E-ANEW (Warriner et al., 2013) to provide the word-level VA annotation. For sentence-level CL, we adopted a corpus from the training split of EmoBank (Buechel and Hahn, 2016, 2017), which contains 40,000 sentences with valence annotation.

The empirical experiments were conducted on four sentiment-related tasks: phrase-level intensity prediction, sentence-level classification and regression, and aspect-level sentiment analysis. The details are presented as follows.

Phrase-level Intensity Prediction. The dataset was the SemEval-2016 task 7 (general English sub-task) (Kiritchenko et al., 2016), which contained 2,799 terms, including 1,330 **words** and 1,469 **phrases**. Each term was annotated within the range of [0, 1]. The metrics contain Kendall’s rank correlation coefficient (k) and Spearman’s rank correlation coefficient (ρ).

Sentence-level Classification. The benchmarks include Stanford Sentiment Treebank (**SST-2/5**) (Socher et al., 2013), Movie Review (**MR**) (Pang and Lee, 2005), and Customer Review (**CR**) (Conneau and Kiela, 2018), **IMDB** (Maas et al., 2011) and **Yelp-2/5** (Zhang et al., 2013). The evaluation metric is accuracy.

Sentence-level Regression. Instead of using binary or fine-grained labels, both the test split of **Emobank** (Buechel and Hahn, 2016, 2017) and **Facebook** (Preotiuc-Pietro et al., 2016) use real-valued VA ratings in the range of (1, 9). The valence denotes the degree of positive and negative sentiment. The metrics are mean absolute error (MAE) and Pearson’s correlation coefficient (r).

Aspect-level Sentiment Classification. This task requires the identification of aspects in text, as well as the assignment of sentiment labels to those aspects. The datasets are SemEval-2014 task 4 (Pontiki et al., 2014) in laptop (**Lap14**) and restaurant (**Res14**). The evaluation metric is accuracy and F_1 -score.

Table 1 summarizes the details of these datasets, including the amount of training, validation, and test splits. For MR, IMDB, and Yelp-2/5, a subset was randomly selected from the training set for validation. For CR, Emobank, and Facebook, the evaluation was performed using a 5-fold cross-validation. For each round, each dataset was randomly split into a training, development, and test set with a 6:2:2 ratio.

4.2. Baselines

The proposed SoftMCL was compared with several pre-trained models, pre-trained by general tasks, i.e., MLM and NSP, and sentiment-related tasks.

The general-purpose pre-trained models include BERT (Devlin et al., 2019), XLNet (Yang et al., 2019), RoBERTa (Liu et al., 2019), and DeBERTa (He et al., 2021). The sentiment-aware pretrained models include BERT-PT (Xu et al., 2019), SentiBERT (Yin et al., 2020), SentiLARE (Ke et al., 2020), SENTIX (Zhou et al., 2020), and SCAPT (Li et al., 2021).

For phrase-level intensity prediction, we further report the empirical results from the original study with the same implementation settings (Yu et al., 2018a,b).

4.3. Implementation Details

The backbone model we used is DeBERTa (He et al., 2021), which is initialized from the pre-trained checkpoint of the base version using the HuggingFace toolkit (Wolf et al., 2020). We use the AdamW optimizer (Kingma and Ba, 2015) with a linear learning rate scheduler with 10% warm-up steps for pre-training. The learning rate is initialized with $2e-5$. The maximum length is truncated to 128. The batch size of contrastive learning is 64. The steps of warm-up and pre-training are 1,500 and 20,000, respectively.

For fine-tuning, the hyper-parameters were selected by using a grid search strategy. The number of epochs for each model was set depending on

Model	Sentence-level sentiment classification							Aspect-level sentiment classification			
	SST-2	SST-5	MR	CR	IMDB	Yelp-2	Yelp-5	Lap14		Rest14	
	Acc ↑	Acc ↑	Acc ↑	Acc ↑	Acc ↑	Acc ↑	Acc	Acc ↑	F_1 -score ↑	Acc ↑	F_1 -score ↑
General pretrained models											
BERT	90.82	53.37	87.52	88.22	93.87	97.74	70.16	78.53	73.11	83.77	76.06
XLNet	92.04	56.33	89.45	88.33	96.21	97.41	70.23	80.00	75.88	84.93	76.70
RoBERTa	92.28	54.89	89.41	88.47	94.68	97.98	70.12	81.03	77.16	86.07	79.21
DeBERTa	94.82	56.89	89.52	88.49	94.92	97.88	70.22	81.14	77.22	86.12	79.43
Sentiment-aware pre-trained models											
BERT-PT	91.42	53.24	87.30	86.52	93.99	97.77	69.90	78.46	73.82	85.58	77.99
SentiBERT	92.18	56.87	88.59	88.24	94.04	97.66	69.94	76.87	71.74	83.71	75.42
SentiLARE	94.28	58.59	90.82	89.52	95.71	98.22	71.57	82.16	78.70	88.32	81.63
SENTIX	93.30	55.57	-	88.72	94.78	97.83	-	80.56	-	87.32	-
SCAPT	-	-	-	-	-	-	-	82.76	79.15	89.11	83.91
Proposed model with ablation study											
SoftMCL	96.54	59.82	92.56	90.02	96.85	98.29	71.76	83.82	81.08	89.23	83.95
w/o WP	94.01	58.25	90.14	87.67	94.31	95.73	69.88	81.63	78.96	86.91	81.75
w/o SP	93.65	58.04	89.79	87.32	93.96	95.35	69.62	81.31	78.66	86.57	81.43
w/o MoCL	91.33	56.61	87.55	85.15	91.62	92.99	67.89	79.29	76.7	84.42	79.41

Table 3: Comparative results of different models on sentence-level and aspect-level sentiment classification tasks.

Model	EmoBank		Facebook	
	MAE ↓	r ↑	MAE ↓	r ↑
General pre-trained models				
BERT	0.581	0.521	0.719	0.635
XLNet	0.523	0.589	0.706	0.654
RoBERTa	0.518	0.592	0.694	0.642
DeBERTa	0.514	0.591	0.711	0.631
Sentiment-aware pre-trained models				
BERT-PT	0.506	0.610	0.708	0.631
SentiBERT	0.505	0.612	0.688	0.677
SentiLARE	0.498	0.615	0.669	0.671
Proposed model with ablation study				
SoftMCL	0.462	0.639	0.642	0.685
w/o WP	0.475	0.629	0.663	0.661
w/o SP	0.495	0.620	0.734	0.634
w/o MoCL	0.492	0.622	0.710	0.662

Table 4: Comparative results of different models on sentence-level regression task to predict valence ratings.

an early strategy, where the patient was set to 3. The model exits when the loss does not decrease in three epochs.

4.4. Comparative Results

This section compared the proposed SoftMCL against several previous methods on four sentiment-related tasks. Tables 2, 3, and 4 summarize the empirical results of the proposed SoftMCL against different baselines on four different tasks. As indicated, the proposed SoftMCL consistently outperformed all baselines for different tasks, demonstrating the superiority of the proposed SoftMCL in capturing affective information in addition to semantic learning.

Moreover, the proposed SoftMCL performs sig-

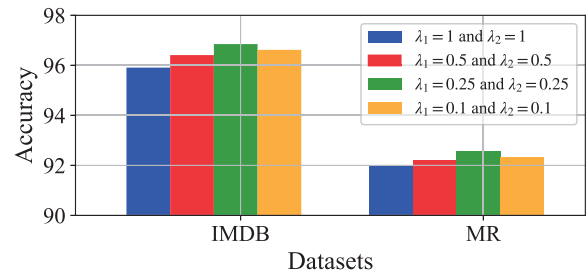


Figure 3: The effect of different balance coefficient.

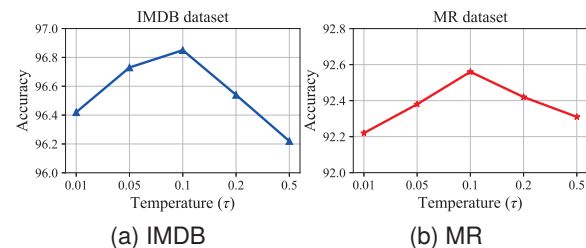


Figure 4: The effect of different temperature.

nificantly better than DeBERTa on all datasets for four different sentiment-related tasks, indicating the effectiveness of the proposed momentum contrastive learning method. Compared with other sentiment-aware pre-trained methods, the improvement mainly comes from twofolds. First, the valence ratings are introduced as soft labels to supervise contrastive learning. Further, the CL is performed on both word- and sentence-level to guide the model to learn more fine-grained sentiment features.

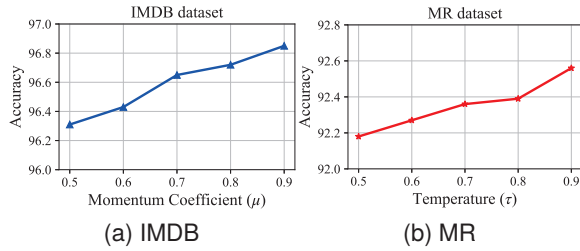


Figure 5: The effect of different momentum coefficient.

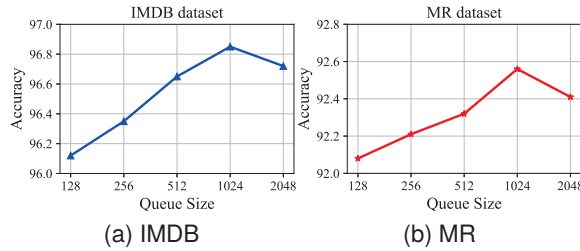


Figure 6: The effect of different queue size.

4.5. Ablation Study

We successively removed word-level pre-training (WP), sentence-level pre-training (SP), and the momentum CL (MoCL) to investigate whether it would reduce performance. The bottom rows of Tables 2, 3, and 4 show the ablation results, indicating each removal produces a varying degree of performance decline. This also proves the necessity of tailoring pre-training paradigms for sentiment-related tasks. The word- and sentence-level pre-training tasks capture affective information in different granularity, and combining multi-granularity pre-training benefits performance improvement.

Further, SoftMCL w/o sentence-level pre-training is worse than w/o word-level pre-training in most circumstances. The performance decline can be consistently observed on aspect-level sentiment classification tasks. The rationale is that the global context is essential for analyzing aspect sentiments while focusing only on words leads to less robust prediction.

By removing the momentum queue and encoder, the proposed SoftMCL will degenerate to only using contrastive samples from the same batch for CL. This will decrease the number of contrastive samples, thus limiting the performance of the proposed SoftMCL.

4.6. Impacts of Parameters

Balance Coefficient. The hyperparameters λ_1 and λ_2 in Eq. (8) were applied to control the contribution of word- and sentence-level CL to the total training. Figure 3 investigates the effect of different

combinations of these two parameters. The observation is that the MLM loss is still essential for the pre-training. When $\lambda_1 = 0.25$ and $\lambda_2 = 0.25$, the SoftMCL can achieve the best results on IMDB and MR datasets.

Temperature. The temperature τ in the training loss of CL controls the density of the softmax function. Figure 4 shows the effect of temperature on IMDB and MR of the SoftMCL. As indicated, all the best results can be achieved using a temperature $\tau = 0.1$. Smaller temperatures benefit training more than higher ones, but extremely low temperatures are more challenging for CL due to numerical instability.

Momentum Coefficient. The momentum coefficient μ can be used to update the parameters of the momentum encoder, making the momentum encoder evolve more smoothly than the original encoder. Figure 5 shows the effect of the momentum coefficient on IMDB and MR datasets. As indicated, a relatively large momentum, e.g., $\mu = 0.9$ works much better than a smaller one, suggesting that a slowly evolving original encoder is essential to utilize the momentum queue.

Queue Size. The size of the momentum queue determines the number of contrastive samples that can be stored. Figure 6 evaluates the impact of queue size on the performance of the SoftMCL. As indicated, the best performance can be achieved when the size is 1024. A large queue requires better hardware but introduces more noise samples.

4.7. Training Analysis

Figure 7 shows the training loss curve of the proposed SoftMCL on word- and sentence-level pre-training. The SoftMCL model gradually converges within 40,000 steps for word-level pre-training since enough emotional tokens can be introduced from each batch for word-level CL training. Even if the batch size setting for sentence-level pre-training limits the number of samples used for CL in a batch, the SoftMCL can still converge quickly within 20,000 steps. This is mainly determined by the contrastive samples introduced by the momentum queue. The soft samples can be stored and applied for training in other batches. The higher the quality of the contrastive samples the momentum queue provided, the better the representation can be learned by the SoftMCL.

5. Conclusion

This study proposed a soft momentum contrastive learning for sentiment-aware pre-training on both the word- and sentence-level to enhance the PLM’s ability to learn affective information. The continuous affective ratings are adopted to guide the training

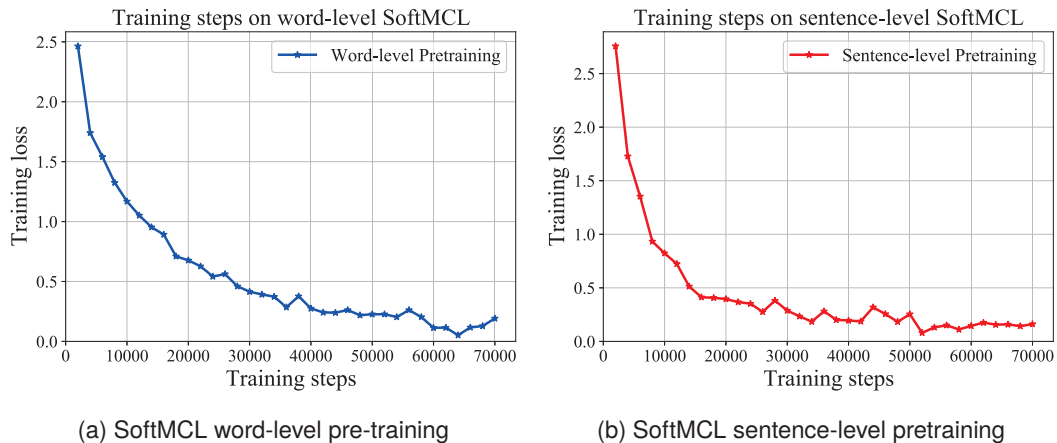


Figure 7: Training curve of the proposed SoftMCL on word- and sentence-level pretraining.

of CL precisely. A momentum queue was introduced to expand the contrastive samples, allowing storing and involving more negatives to overcome the limitations of hardware platforms. Extensive experiments on four different sentiment-related tasks have proven the effectiveness of the proposed SoftMCL method.

Future works will attempt to incorporate the sentiment information into decoder-only architecture to achieve the sentiment expression ability for generative models.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant Nos. 61966038 and 62266051, the Ministry of Science and Technology (MOST), Taiwan, ROC, under Grant No. MOST110-2628-E-155-002. The authors would like to thank the anonymous reviewers for their constructive comments.

Bibliographical References

Mohamed Abdalla, Magnus Sahlgren, and Graeme Hirst. 2019. Enriching Word Embeddings with a Regressor Instead of Labeled Corpora. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-2019)*, pages 6188–6195.

Ameeta Agrawal, Aijun An, and Manos Papagelis. 2018. Learning Emotion-enriched Word Representations. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING-2018)*, pages 950–961.

Sven Buechel and Udo Hahn. 2016. Readers vs. writers vs. texts : Coping with different perspec-

tives of text understanding in emotion annotation. In *Proceedings of the 11th Linguistic Annotation Workshop in EACL 2017*, pages 1–12.

Sven Buechel and Udo Hahn. 2017. EmoBank: Studying the Impact of Annotation Perspective and Representation Format on Dimensional Emotion Analysis. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL-2017)*, pages 578–585.

Rafael A. Calvo and Sunghwan Mac Kim. 2013. Emotions in Text: Dimensional and Categorical Models. *Computational Intelligence*, 29(3):527–543.

Changyou Chen, Jiali Duan, Yiran Chen, Jianyi Zhang, Son Dinh Tran, Yi Xu, Liqun Chen, Belinda Zeng, and Trishul Chilimbi. 2022. Why do We Need Large Batchsizes in Contrastive Learning? A Gradient-Bias Perspective. In *Advances in Neural Information Processing Systems*, volume 35, pages 1–16.

Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing. In *Proceedings of the 25th International Conference on Machine Learning (ICML-2008)*, pages 160–167.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (Almost) from scratch. *Journal of Machine Learning Research*, 12:2493–2537.

Alexis Conneau and Douwe Kiela. 2018. SentEval: An Evaluation Toolkit for Universal Sentence Representations. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC-2018)*, pages 1699–1704.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL/HLT-2019)*, pages 4171–4186.
- Paul Ekman. 1992. An Argument for Basic Emotions. *Cognition & Emotion*, 6(3):169–200.
- Shuai Fan, Chen Lin, Haonan Li, Zhenghao Lin, Jinsong Su, Hang Zhang, Yeyun Gong, Jian Guo, and Nan Duan. 2022. Sentiment-aware Word and Sentence Level Pre-training for Sentiment Analysis.
- Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP-2017)*, pages 1615–1625.
- Peng Fu, Zheng Lin, Fengcheng Yuan, Weiping Wang, and Dan Meng. 2018. Learning Sentiment-specific Word Embedding via Global Sentiment Representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI-2018)*, pages 4808–4815.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP-2021)*, pages 6894–6910.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR-2020)*, pages 9726–9735. IEEE.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing. In *International Conference on Learning Representations (ICLR-2021)*, pages 1–16.
- Nora Kassner and Hinrich Schütze. 2020. Negated and Misprimed Probes for Pretrained Language Models: Birds Can Talk, But Cannot Fly. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL-2020)*, pages 7811–7818.
- Pei Ke, Haozhe Ji, Siyang Liu, Xiaoyan Zhu, and Minlie Huang. 2020. SentiLARE: Sentiment-aware Language Representation Learning with Linguistic Knowledge. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP-2020)*, pages 6975–6988, Online. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the International Conference on Learning Representation (ICLR-2015)*, pages 1–15.
- Svetlana Kiritchenko, Saif M Mohammad, and Mohammad Salameh. 2016. SemEval-2016 Task 7: Determining sentiment intensity of English and Arabic phrases. In *Proceedings of International Workshop on Semantic Evaluation (SemEval-2016)*, pages 42–51.
- Man Lan, Zihua Zhang, Yue Lu, and Ju Wu. 2016. Three convolutional neural network-based models for learning sentiment word vectors towards sentiment analysis. In *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN-2016)*, pages 3172–3179.
- Lung Hao Lee, Jian Hong Li, and Liang Chih Yu. 2022. Chinese EmoBank: Building Valence-Arousal Resources for Dimensional Sentiment Analysis. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 21(4):1–18.
- Zhengyan Li, Yicheng Zou, Chong Zhang, Qi Zhang, and Zhongyu Wei. 2021. Learning Implicit Sentiment in Aspect-based Sentiment Analysis with Supervised Contrastive Pre-training. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Proceedings (EMNLP-2021)*, pages 246–256.
- 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 *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (ACL-2011)*, pages 142–150.
- Bo Pang and Lillian Lee. 2005. Seeing Stars: Exploiting Class Relationships for Sentiment Categorization with Respect to Rating Scales. In

- Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics (ACL-2005)*, pages 115–124.
- Robert Plutchik. 1991. *The emotions*, university edition.
- 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, Proceedings (SemEval-2014)*, pages 27–35.
- Daniel Preotiuc-Pietro, H. Andrew Schwartz, Gregory Park, Johannes C. Eichstaedt, Margaret Kern, Lyle Ungar, and Elizabeth P. Shulman. 2016. *Modeling Valence and Arousal in Facebook Posts*. In *Proceedings of Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA-2016)*, pages 9–15.
- Yafeng Ren, Yue Zhang, Meishan Zhang, and Donghong Ji. 2016. Improving Twitter sentiment classification using topic-enriched multi-prototype word embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-2016)*, pages 3038–3044.
- James A Russell. 1980. A Circumplex Model of Affect. *Journal of Personality and Social Psychology*, 39(6):1161–1178.
- Richard Socher, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive Deep Models for Semantic Compositionality over a Sentiment Treebank. In *EMNLP 2013 - 2013 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pages 1631–1642.
- Duyu Tang, Furu Wei, Bing Qin, Nan Yang, Ting Liu, and Ming Zhou. 2016. Sentiment embeddings with applications to sentiment analysis. *IEEE Transactions on Knowledge and Data Engineering*, 28(2):496–509.
- Hao Tian, Can Gao, Xinyan Xiao, Hao Liu, Bolei He, Hua Wu, Haifeng Wang, and Feng Wu. 2020. SKEP: Sentiment Knowledge Enhanced Pre-training for Sentiment Analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL-2020)*, pages 4067–4076.
- Akira Utsumi. 2019. Refining Pretrained Word Embeddings Using Layer-wise Relevance Propagation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP-2018)*, pages 4840–4846.
- Jin Wang, Liang-Chih Yu, K. Robert Lai, and Xuejie Zhang. 2019. Investigating Dynamic Routing in Tree-Structured LSTM for Sentiment Analysis. 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-2019)*, pages 3423–3428.
- Jin Wang, Liang Chih Yu, K. Robert Lai, and Xuejie Zhang. 2020. Tree-structured Regional CNN-LSTM Model for Dimensional Sentiment Analysis. *IEEE/ACM Transactions on Audio Speech and Language Processing*, 28:581–591.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of Valence, Arousal, and Dominance for 13,915 English Lemmas. *Behavior Research Methods*, 45(4):1191–1207.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick Von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations (EMNLP-2020)*, pages 38–45.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. XLNet: Generalized Autoregressive Pre-training for Language Understanding. In *Advances in Neural Information Processing Systems (NeurIPS-2019)*, volume 32, pages 5753–5763. Curran Associates, Inc.
- Zhe Ye, Fang Li, and Timothy Baldwin. 2018. Encoding Sentiment Information into Word Vectors for Sentiment Analysis. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING-2018)*, pages 997–1007.
- 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 (ACL-2020)*, pages 3695–3706.
- Liang-Chih Yu, Jin Wang, K. Robert Lai, and Xuejie Zhang. 2018a. Refining Word Embeddings for Sentiment Analysis. In *Proceedings of the*

2017 Conference on Empirical Methods in Natural Language Processing (EMNLP-2017), pages 534–539.

Liang-Chih Yu, Jin Wang, K. Robert Lai, and Xuejie Zhang. 2018b. Refining Word Embeddings Using Intensity Scores for Sentiment Analysis. *IEEE/ACM Transactions on Audio Speech and Language Processing*, 26(3):671–681.

Xiang Zhang, Junbo Zhao, and Yann LeCun. 2013. Character-level Convolutional Networks for Text Classification. In *Proceedings of Advances in Neural Information Processing Systems (NIPS-2015)*, pages 3057–3061.

Jie Zhou, Junfeng Tian, Rui Wang, Yuanbin Wu, Wenming Xiao, and Liang He. 2020. SENTIX: A Sentiment-aware Pre-trained Model for Cross-domain Sentiment Analysis. In *Proceedings of the 28th International Conference on Computational Linguistics (COLING-2020)*, pages 568–579.

A. Related Works

The appendix briefly reviewed the related works of sentiment embeddings and sentiment-aware pre-training.

A.1. Sentiment-enhanced Embeddings

Using distributed representation to map words into latent space, the technique of word embeddings could make semantically related words appear closer to one another. These word embeddings can serve as the input for various neural network models to leverage the contextual information in large corpora, thus benefiting many NLP tasks (Wang et al., 2019, 2020). Unfortunately, the traditional semantic word embeddings are not sufficiently practical or feasible when directly used for sentiment analysis tasks. Several studies (Agrawal et al., 2018) have reported that good is very similar to bad in vector space because these words are commonly used in similar contexts despite being opposite in sentiment. As a result, two sentences with similar contexts, respectively containing good and bad, will be assigned a similar sentence representation.

Previous studies have suggested using sentiment-aware pre-training (Abdalla et al., 2019; Fu et al., 2018; Lan et al., 2016; Ren et al., 2016; Tang et al., 2016) or sentiment refinement (Utsumi, 2019; Yu et al., 2018b) to encode external knowledge into word representation to enrich the model with sentiment information.

Sentiment-aware Pre-training. The C&W model (Collobert and Weston, 2008; Collobert et al., 2011)

is typically extended to train on a corpus with sentiment labels to learn sentiment embeddings from scratch. Tang et al. (2016) proposed an end-to-end architecture with different loss functions to incorporate sentiment and semantic information of words. Ren et al. (2016) further integrated topic information into sentiment embeddings to learn a multi-prototype topic and sentiment-enriched word embeddings. Felbo et al. (2017) used a bi-directional LSTM model trained from scratch on almost one billion tweets to obtain context-aware sentiment embeddings.

Sentiment Refinement. An alternative method is to adjust pre-trained word embeddings with external word- or sentence-level knowledge using an embedding refinement model and then leverage the modified embeddings for sentiment analysis.

The refinement model could use affective lexicons to enrich word-level sentiment information to pre-trained word embeddings. Yu et al. (2018a) and (2018b) extracted VA ratings from affective lexicons to pull similar sentiment words together in the latent space. Ye et al. (2018) used a joint learning approach to integrate sentiment into word embeddings by predicting both word- and sentence-level sentiment scores.

A.2. Sentiment-aware Transformers

For sentiment-related tasks, several studies proposed to train a transformer-based model by being supervised with sentiment labels as either a pre-training or a post-training task. For example, ? proposed BERT-PT, which conducts sentiment-related post-training on the corpora to benefit aspect-level sentiment analysis. Tian et al. (2020) proposed SKEP to introduce three predictions with sentiment knowledge to learn a unified sentiment representation for various sentiment applications. Zhou et al. (2020) proposed SENTIX to extract domain-invariant sentiment features from large-scale review datasets, which can then perform cross-domain sentiment classification without fine-tuning. (Yin et al., 2020) proposed SentiBERT, which applied a bi-attention mechanism on top of the BERT representation to learn phrase-level compositional semantics. Ke et al. (2020) proposed SentiLARE, which adopted a label-aware masked language modeling to learn sentiment-aware language representation. (Li et al., 2021) proposed SCAPT to obtain implicit and explicit sentiment polarity by aligning the representation of implicit sentiment expressions to those with the same polarity.