

SentiBERT: A Transferable Transformer-Based Architecture for Compositional Sentiment Semantics

Da Yin[♣], Tao Meng[♣], Kai-Wei Chang[♣]

[♣] Peking University

[♣] University of California, Los Angeles

wade_yin9712@pku.edu.cn, tmeng@cs.ucla.edu, kw@kwchang.net

Abstract

We propose *SentiBERT*, a variant of BERT that effectively captures compositional sentiment semantics. The model incorporates contextualized representation with binary constituency parse tree to capture semantic composition. Comprehensive experiments demonstrate that *SentiBERT* achieves competitive performance on phrase-level sentiment classification. We further demonstrate that the sentiment composition learned from the phrase-level annotations on SST can be transferred to other sentiment analysis tasks as well as related tasks, such as emotion classification tasks. Moreover, we conduct ablation studies and design visualization methods to understand *SentiBERT*. We show that *SentiBERT* is better than baseline approaches in capturing negation and the contrastive relation and model the compositional sentiment semantics.

1 Introduction

Sentiment analysis is an important language processing task (Pang et al., 2002, 2008; Liu, 2012). One of the key challenges in sentiment analysis is to model compositional sentiment semantics. Take the sentence “*Frenetic but not really funny.*” in Figure 1 as an example. The two parts of the sentence are connected by “*but*”, which reveals the change of sentiment. Besides, the word “*not*” changes the sentiment of “*really funny*”. These types of negation and contrast are often difficult to handle when the sentences are complex (Socher et al., 2013; Tay et al., 2018; Xu et al., 2019).

In general, the sentiment of an expression is determined by the meaning of tokens and phrases and the way how they are syntactically combined. Prior studies consider explicitly modeling compositional sentiment semantics over constituency structure with recursive neural networks (Socher et al., 2012,




Figure 1: Illustration of the challenges of learning sentiment semantic compositionality. The blue nodes represent token nodes. The colors of phrase nodes in the binary constituency tree represent the sentiment of phrases. The red boxes show that the sentiment changes from the child node to the parent node due to negation and contrast.

2013). However, these models that generate representation of a parent node by aggregating the local information from child nodes, overlook the rich association in context.

In this paper, we propose *SentiBERT* to incorporate recently developed contextualized representation models (Devlin et al., 2019; Liu et al., 2019) with the recursive constituency tree structure to better capture compositional sentiment semantics. Specifically, we build a simple yet effective attention network for composing sentiment semantics on top of BERT (Devlin et al., 2019). During training, we follow BERT to capture contextual information by masked language modeling. In addition, we instruct the model to learn composition of meaning by predicting sentiment labels of the phrase nodes.

Results on phrase-level sentiment classification on Stanford Sentiment Treebank (SST) (Socher et al., 2013) indicate that *SentiBERT* improves significantly over recursive networks and the base-

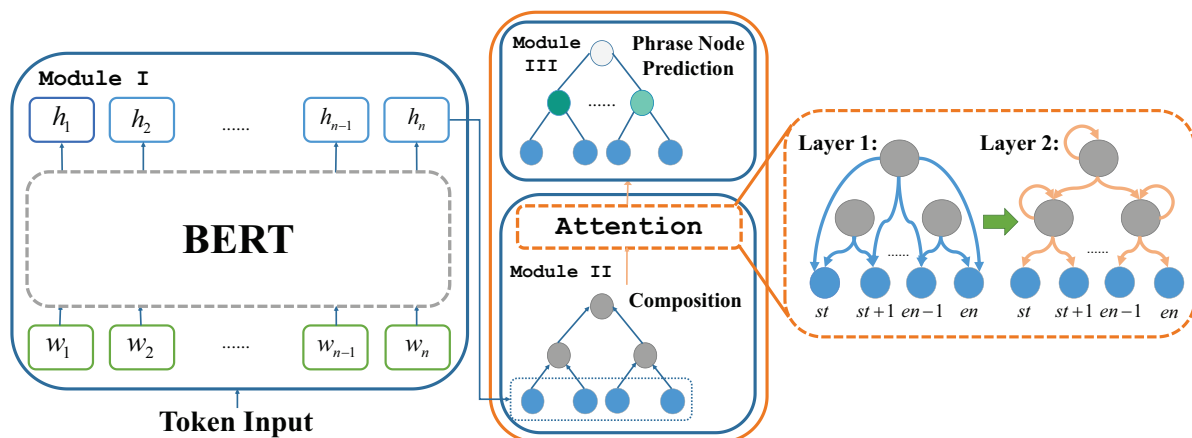


Figure 2: The architecture of SentiBERT. Module I is the BERT encoder; Module II denotes the semantic composition module based on an attention mechanism; Module III is a predictor for phrase-level sentiment. The semantic composition module is a two layer attention-based network (see Section 3.1) The first layer (**Attention to Tokens**) generates representation for each phrase based on the token it covers and the second layer (**Attention to Children**) refines the phrase representation obtained from the first layer based on its children.

line BERT model. As phrase-level sentiment labels are expensive to obtain, we further explore if the compositional sentiment semantics learned from one task can be transferred to others. In particular, we find that SentiBERT trained on SST can be transferred well to other related tasks such as twitter sentiment analysis (Rosenthal et al., 2017) and emotion intensity classification (Mohammad et al., 2018) and contextual emotion detection (Chatterjee et al., 2019). Furthermore, we conduct comprehensive quantitative and qualitative analyses to evaluate the effectiveness of SentiBERT under various situations and to demonstrate the semantic compositionality captured by the model. The source code is available at <https://github.com/WadeYin9712/SentiBERT>.

2 Related Work

Sentiment Analysis Various approaches have been applied to build a sentiment classifier, including feature-based methods (Hu and Liu, 2004; Pang and Lee, 2004), recursive neural networks (Socher et al., 2012, 2013; Tai et al., 2015), convolution neural networks (Kim, 2014) and recurrent neural networks (Liu et al., 2015). Recently, pre-trained language models such as ELMo (Peters et al., 2018), BERT (Devlin et al., 2019) and SentiLR (Ke et al., 2019) achieve high performance in sentiment analysis by constructing contextualized representation. Inspired by these prior studies, we design a transformer-based neural network model to capture compositional sentence semantics by

leveraging binary constituency parse tree.

Semantic Compositionality Semantic composition (Pelletier, 1994) has been widely studied in NLP literature. For example, Mitchell and Lapata (2008) introduce operations such as addition or element-wise product to model compositional semantics. The idea of modeling semantic composition is applied to various areas such as sentiment analysis (Socher et al., 2013; Zhu et al., 2016), semantic relatedness (Marelli et al., 2014) and capturing sememe knowledge (Qi et al., 2019). In this paper, we demonstrate that the syntactic structure can be combined with contextualized representation such that the semantic compositionality can be better captured. Our approach resembles to a few recent attempts (Harer et al., 2019; Wang et al., 2019) to integrate tree structures into self-attention. However, our design is specific for the semantic composition in sentiment analysis.

3 Model

We introduce SentiBERT, a model that captures compositional sentiment semantics based on constituency structures of sentences. SentiBERT consists of three modules: 1) BERT; 2) a semantic composition module based on an attention network; 3) phrase and sentence sentiment predictors. The three modules are illustrated in Figure 2 and we provide an overview in below.

BERT We incorporate BERT (Devlin et al., 2019) as the backbone to generate contextualized

representation of input sentence.

Semantic Composition Module This module aims to obtain effective phrase representation guided by the contextualized representation and constituency parsing tree. To refine phrase representation based on the structural information and its constituencies, we design a two-level attention mechanism: 1) *Attention to Tokens* and 2) *Attention to Children*.

Phrase Node Prediction `SentiBERT` is supervised by phrase-level sentiment labels. We use cross-entropy as the loss function for learning the sentiment predictor.

3.1 Attention Networks for Sentiment Semantic Composition

In this section, we describe the attention networks for sentiment semantic composition in detail.

We first introduce the notations. $s = [w_1, w_2, \dots, w_n]$ denotes a sentence which consists of n words. $\text{phr} = [\text{phr}_1, \text{phr}_2, \dots, \text{phr}_m]$ denotes the phrases on the binary constituency tree of sentence s . $\mathbf{h} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n]$ is the contextualized representation of tokens after forwarding to a fully-connected layer, where $\mathbf{h}_t \in \mathbb{R}^d$. Suppose st_i and en_i are beginning and end indices of the i -th phrase where $w_{st_i}, w_{st_i+1}, \dots, w_{en_i}$ are constituent tokens of the i -th phrase. The corresponding token representation is $[\mathbf{h}_{st_i}, \mathbf{h}_{st_i+1}, \dots, \mathbf{h}_{en_i}]$. \mathbf{p}_i is the phrase representation of the i -th phrase.

Attention to Tokens Given the contextualized representations of the tokens covered by a phrase. We first generate phrase representation \mathbf{v}_i for a phrase i by the following attention network.

$$\begin{aligned} \mathbf{q}_i &= \frac{1}{en_i - st_i + 1} \sum_{j=st_i}^{en_i} \mathbf{h}_j, \\ t_j &= \text{Attention}(\mathbf{q}_i, \mathbf{h}_j), st_i \leq j \leq en_i, \\ a_j &= \frac{\exp(t_j)}{\sum_{k=st_i}^{en_i} \exp(t_k)}, \\ \mathbf{o}_i &= \sum_{j=st_i}^{en_i} a_j \cdot \mathbf{h}_j. \end{aligned} \quad (1)$$

In Eq. (1), we first treat the averaged representation for each token as the query, and then allocate attention weights according to the correlation with each token. a_j represents the weight distributed to the j -th token. We concatenate the weighted sum \mathbf{o}_i and \mathbf{q}_i and feed it to forward networks.

Lastly, we obtain the initial representation for the phrase $\mathbf{v}_i \in \mathbb{R}^d$ based on the representation of constituent tokens. The detailed computation of attention mechanism is shown in Appendix A.1.

Attention to Children Furthermore, we refine phrase representations in the second layer based on constituency parsing tree and the representations obtained in the first layer. To aggregate information based on hierarchical structure, we develop the following network. For each phrase, the attention network computes correlation with its children in the binary constituency parse tree and itself.

Assume that the indices of child nodes of the i -th phrase are l_{son} and r_{son} . Their representations generated from the first layer are $\mathbf{v}_i, \mathbf{v}_{l_{son}}$, and $\mathbf{v}_{r_{son}}$, respectively. We generate the attention weights $r_{l_{son}}, r_{r_{son}}$ and r_i over the i -th phrase and its left and right children by the following.

$$\begin{aligned} c_{l_{son}} &= \text{Attention}(\mathbf{v}_i, \mathbf{v}_{l_{son}}), \\ c_{r_{son}} &= \text{Attention}(\mathbf{v}_i, \mathbf{v}_{r_{son}}), \\ c_i &= \text{Attention}(\mathbf{v}_i, \mathbf{v}_i), \\ r_{l_{son}}, r_{r_{son}}, r_i &= \text{Softmax}(c_{l_{son}}, c_{r_{son}}, c_i). \end{aligned} \quad (2)$$

Then the refined representation of phrase i is computed by

$$\mathbf{f}_i = r_{l_{son}} \cdot \mathbf{v}_{l_{son}} + r_{r_{son}} \cdot \mathbf{v}_{r_{son}} + r_i \cdot \mathbf{v}_i.$$

Finally, we concatenate the weighted sum \mathbf{f}_i and \mathbf{v}_i and feed it to forward networks with SeLU (Klambauer et al., 2017) and GeLU activations (Hendrycks and Gimpel, 2017) and layer normalization (Ba et al., 2016), similar to Joshi et al. (2020) to generate the final phrase representation $\mathbf{p}_i \in \mathbb{R}^d$. Note that when the child of i -th phrase is token node, the attention mechanism will attend to the representation of all the subtokens the token node covers.

3.2 Training Objective of SentiBERT

Inspired by BERT, the training objective of `SentiBERT` consists of two parts: 1) Masked Language Modeling. Some texts are masked and the model learn to predict them. This objective allows the model learn to capture the contextual information as in the original BERT model. 2) Phrase Node Prediction. We further consider training the model to predict the phrase-level sentiment label based on the aforementioned phrase representations. This allows `SentiBERT` learn to capture the compositional sentiment semantics. Similar to BERT, in the

transfer learning setting, pre-trained `SentiBERT` model can be used to initialize the model parameters of a downstream model.

4 Experiments

We evaluate `SentiBERT` on the SST dataset. We then evaluate `SentiBERT` in a transfer learning setting and demonstrate that the compositional sentiment semantics learned on SST can be transferred to other related tasks.

4.1 Experimental Settings

We evaluate how effective `SentiBERT` captures the compositional sentiment semantics on SST dataset (Socher et al., 2013).

The SST dataset has several variants.

- SST-phrase is a 5-class classification task that requires to predict the sentiment of all phrases on a binary constituency tree. Different from Socher et al. (2013), we test the model only on phrases (non-terminal constituents) and ignore its performance on tokens.
- SST-5 is a 5-class sentiment classification task that aims at predicting the sentiment of a sentence. We use it to test if `SentiBERT` learns a better sentence representation through capturing compositional sentiment semantics.
- Similar to SST-5, SST-2 and SST-3 are 2-class and 3-class sentiment classification tasks. However, the granularity of the sentiment classes is different.

Besides, to test the transferability of `SentiBERT`, we consider several related datasets, including Twitter Sentiment Analysis (Rosenthal et al., 2017), Emotion Intensity Classification (Mohammad et al., 2018) and Contextual Emotion Detection (EmoContext) (Chatterjee et al., 2019). Details are shown in Appendix A.2.

We build `SentiBERT` on the HuggingFace library¹ and initialize the model parameters using pre-trained BERT-base and RoBERTa-base models whose maximum length is 128, layer number is 12, and embedding dimension is 768. For the training on SST-phrase, the learning rate is 2×10^{-5} , batch size is 32 and the number of training epochs is 3. For masking mechanism, to put emphasis on

modeling sentiments, the probability of masking opinion words which can be retrieved from SentiWordNet (Baccianella et al., 2010) is set to 20%, and for the other words, the probability is 15%. For fine-tuning on downstream tasks, the learning rate is $\{1 \times 10^{-5} - 1 \times 10^{-4}\}$, batch size is $\{16, 32\}$ and the number of training epochs is 1–5. We use Stanford CoreNLP API (Manning et al., 2014) to obtain binary constituency trees for the sentences of these tasks to keep consistent with the settings on SST-phrase. Note that when fine-tuning on sentence-level sentiment and emotion classification tasks, the objective is to correctly label the root of tree, instead of targeting at the [CLS] token representation as in the original BERT.

4.2 Effectiveness of SentiBERT

We first compare the proposed attention networks (`SentiBERT` w/o BERT) with the following baseline models trained on SST-phrase corpus to evaluate the effectiveness of the architecture design: 1) Recursive NN (Socher et al., 2013); 2) GCN (Kipf and Welling, 2017); 3) Tree-LSTM (Tai et al., 2015); 4) BiLSTM (Hochreiter and Schmidhuber, 1997) w/ Tree-LSTM. To further understand the effect of using contextualized representation, we compare `SentiBERT` with the vanilla pre-trained BERT and its variants which combine the four mentioned baselines and BERT. The training settings remain the same with `SentiBERT`. We also initialize `SentiBERT` with pre-trained parameters of RoBERTa (`SentiBERT` w/ RoBERTa) and further compare it with its variants.

As shown in Table 1, `SentiBERT` and `SentiBERT` w/ RoBERTa substantially outperforms their corresponding variants and the networks merely built on the tree. Specifically, we first observe that though our attention network (`SentiBERT` w/o BERT) is simple, it is competitive with Recursive NN, GCN and Tree-LSTM. Besides, `SentiBERT` largely outperforms `SentiBERT` w/o BERT by leveraging contextualized representation. Moreover, the results manifest that `SentiBERT` and `SentiBERT` w/ RoBERTa outperform the BERT and RoBERTa, indicating the importance of incorporating syntactic guidance.

4.3 Transferability of SentiBERT

Though the designed models are effective, we are curious how beneficial the compositional sentiment semantics learned on SST can be transferred to other tasks. We compare `SentiBERT` with pub-

¹<https://github.com/huggingface>

Models	SST-phrase	SST-5
Recursive NN	58.33	46.53
GCN	60.89	49.34
Tree-LSTM	61.71	50.07
BiLSTM w/ Tree-LSTM	61.89	50.45
BERT w/ Mean pooling	64.53	50.68
BERT w/ GCN	65.23	54.56
BERT w/ Tree-LSTM	67.39	55.89
RoBERTa w/ Mean pooling	67.73	56.34
SentiBERT w/o BERT	61.04	50.31
SentiBERT	68.31	56.10
SentiBERT w/ RoBERTa	68.98	56.87

Table 1: The averaged accuracies on SST-phrase and SST-5 tasks (%) for 5 runs. For baselines vanilla BERT and RoBERTa, we use mean-pooling on token representation of top layer to get phrase and sentence representation.

Models	SST-2 (Dev)	SST-3	Twitter
BERT	92.39	73.78	70.0
BERT w/ Mean pooling	92.33	74.35	69.7
XLNet	93.23	75.89	70.7
RoBERTa	94.31	78.04	71.1
SentiBERT w/o BERT	86.57	68.32	64.9
SentiBERT w/o Masking	92.48	76.95	70.7
SentiBERT w/o Pre-training	92.44	76.78	70.8
SentiBERT	92.78	77.11	70.9
SentiBERT w/ RoBERTa	94.72	78.69	71.5

Table 2: The averaged results on sentence-level sentiment classification (%) for 5 runs. For SST-2,3, the metric is accuracy; for Twitter Sentiment Analysis, we use averaged recall value.

lished models BERT, XLNet, RoBERTa and their variants on benchmarks mentioned in Section 4.1. Specifically, ‘BERT’ indicates the model trained on the raw texts of the SST dataset. ‘BERT w/ Mean pooling’ denotes the model trained on SST, whose phrase and sentence representation is computed by mean pooling on tokens. ‘BERT w/ Mean pooling’ merely leverages the phrases’ range information rather than syntactic structural information.

Sentiment Classification Tasks The evaluation results of sentence-level sentiment classification on the three tasks are shown in Table 2. Despite the difference among tasks and datasets, from experimental results, we find that SentiBERT has competitive performance compared with various baselines. SentiBERT achieves higher performance than the vanilla BERT and XLNet in tasks such as SST-3 and Twitter Sentiment Analysis. Besides, SentiBERT significantly outperform

Models	Emotion Intensity	EmoContext
BERT	65.2	73.49
RoBERTa	66.4	74.20
SentiBERT w/o Pre-training	66.0	73.81
SentiBERT	66.5	74.23
SentiBERT w/ RoBERTa	67.2	74.67

Table 3: The averaged results on several emotion classification tasks (%) for 5 runs. For Emotion Intensity Classification task, the metric is averaged Pearson Correlation value of the four subtasks; for EmoContext, we follow the standard metrics used in Chatterjee et al. (2019) and use F1 score as the evaluation metric.

SentiBERT w/o BERT. This demonstrates the importance of leveraging pre-trained BERT model. Moreover, SentiBERT outperforms BERT w/ Mean pooling. This indicates the importance of modeling the compositional structure of sentiment.

Emotion Classification Tasks Emotion detection is different from sentiment classification. However, these two tasks are related. The task aims to classify fine-grained emotions, such as happiness, fearness, anger, sadness, etc. It is challenging compared to sentiment analysis because of various emotion types. We fine-tune SentiBERT and SentiBERT w/ RoBERTa on Emotion Intensity Classification and EmoContext. Table 3 shows the results on the two emotion classification tasks. Similar to the results in sentiment classification tasks, SentiBERT obtains the best results, further justifying the transferability of SentiBERT.

5 Analysis

We conduct experiments on SST-phrase using BERT-base model as backbone to demonstrate the effectiveness and interpretability of the SentiBERT architecture in terms of semantic compositionality. We also explore potential of the model when lacking phrase-level sentiment information. In order to simplify the analysis of the change of sentiment polarity, we convert the 5-class labels to 3-class: the classes ‘very negative’ and ‘negative’ are converted to be ‘negative’; the classes ‘very positive’ and ‘positive’ are converted to be ‘positive’; the class ‘neutral’ remains the same. The details of statistical distribution in this part is shown in Appendix A.3.

We consider the following baselines to evaluate the effectiveness of each component in SentiBERT. First we design BERT w/ Mean pooling as a base model, to demonstrate the ne-




Figure 3: Evaluation for local difficulty. The figure shows the accuracy difference on phrase node sentiment prediction with BERT w/ Mean pooling for different local difficulty.

cessity of incorporating syntactic guidance and implementing aggregation on it. Then we compare `SentiBERT` with alternative aggregation approaches, `Tree-LSTM`, `GCN` and `w/o Attention to Children`.

5.1 Semantic Compositionality

We investigate how effectively `SentiBERT` captures compositional sentiment semantics. We focus on how the representation in `SentiBERT` captures the sentiments when the children and parent in the constituency tree have different sentiments (i.e., sentiment switch) as shown in the red boxes of Figure 1. Here we focus on the sentiment switches between phrases. We assume that the more the sentiment switches, the harder the prediction is.

We analyze the model under the following two scenarios: *local difficulty* and *global difficulty*. Local difficulty is defined as the number of sentiment switches between a phrase and its children. As we consider binary constituency tree. The maximum number of sentiment switches for each phrase is 2. Global difficulty indicates number of sentiment switches in the entire constituency tree. The maximum number of sentiment switches in the test set is 23. The former is a phrase-level analysis and the latter is sentence level.

We compare `SentiBERT` with aforementioned baselines. We group all the nodes and sentences in the test set by local and global difficulty. Results are shown in Figure 3 and Figure 4. Our model achieves better performance than baselines in all situations. Also, we find that with the increase of difficulty, the gap between our models




Figure 4: Evaluation for global difficulty. The figure shows the accuracy difference on phrase node sentiment prediction with BERT w/ Mean pooling for different global difficulty.

and baselines becomes larger. Especially, when the sentiment labels of both children are different from the parent node (i.e., local difficulty is 2), the performance gap between `SentiBERT` and `BERT w/ Tree-LSTM` is about 7% accuracy. It also outperforms the baseline `BERT` model with mean pooling by 15%. This validates the necessity of structural information for semantic composition and the effectiveness of our designed attention networks for leveraging the hierarchical structures.

5.2 Negation and Contrastive Relation

Next, we investigate how `SentiBERT` deals with negations and contrastive relation.

Negation: Since the negation words such as ‘*no*’, ‘*n’t*’ and ‘*not*’ will cause the sentiment switches, the number of negation words also reflects the difficulty of understanding sentence and its constituencies. We first group the sentences by the number of negation words, and then calculate the accuracy of the prediction on their constituencies respectively. In test set, as there are at most six negation words and the amount of sentences with above three negation words is small, we separate all the data into three groups.

Results are provided in Figure 5. We observe `SentiBERT` performs the best among all the models. Similar to the trend in local and global difficulty experiments, the gap between `SentiBERT` and other baselines becomes larger with increase of negation words. The results show the ability of `SentiBERT` when dealing with negations.




Figure 5: Evaluation for negation. We show the accuracy difference with BERT w/ Mean pooling.

Models	Accuracy
BERT w/ Mean pooling	26.1
BERT w/ Tree-LSTM	28.5
BERT w/ GCN	29.4
SentiBERT w/o Attention to Children	29.8
SentiBERT	30.7

Table 4: Evaluation for contrastive relation (%). We show the accuracy for triple-lets (‘X but Y’, ‘X’, ‘Y’). X and Y must be phrases in our experiments.

Contrastive Relation: We evaluate the effectiveness of SentiBERT with regards to tackling contrastive relation problem. Here, we focus on the contrastive conjunction “*but*”. We pick up the sentences containing word ‘but’ of which the sentiments of left and right parts are different. In our analysis, a ‘X but Y’ can be counted as correct if and only if the sentiments of all the phrases in triple-let (‘X but Y’, ‘X’ and ‘Y’) are predicted correctly. Table 4 demonstrates the results. SentiBERT outperforms other variants of BERT about 1%, demonstrating its ability in capturing contrastive relation in sentences.

5.3 Case Study

We showcase several examples to demonstrate how SentiBERT performs sentiment semantic composition. We observe the attention distribution among hierarchical structures. In Figure 7, we demonstrate two sentences of which the sentiments of all the phrases are predicted correctly. We also visualize the attention weights distributed to the child nodes and the phrases themselves to see which part might contribute more to the sentiment of those phrases.

SentiBERT performs well in several aspects.

First, SentiBERT tends to attend to adjectives such as ‘*frenetic*’ and ‘*funny*’, which contribute to the phrases’ sentiment. Secondly, facing negation words, SentiBERT considers them and a switch can be observed between the phrases with and without negation word (e.g., ‘*not really funny*’ and ‘*really funny*’). Moreover, SentiBERT can correctly analyze the sentences expressing different sentiments in different parts. For the first case, the model concentrates more on the part after ‘*but*’.

5.4 Amount of Phrase-level Supervision

We are also interested in analyzing how much phrase-level supervision SentiBERT needs in order to capture the semantic compositionality. We vary the amount of phrase-level annotations used in training SentiBERT. Before training, we randomly sample 0% to 100% with a step of 10% of labels from SST training set. After pre-training on them, we fine-tune SentiBERT on tasks SST-5, SST-3 and Twitter Sentiment Analysis. During fine-tuning, for the tasks which use phrase-level annotation, such as SST-5 and SST-3, we use the same phrase-level annotation during pre-training and the sentence-level annotation; for the tasks which do not have phrase-level annotation, we merely use the sentence-level annotation.

Results in Figure 6 show that with about 30%-50% of the phrase labels on SST-5 and SST-3, the model is able to achieve competitive results compared with XLNet. Even without any phrase-level supervision, using 70%-80% of phrase labels in pre-training allows SentiBERT competitive with XLNet on the Twitter Sentiment Analysis dataset.

Furthermore, we find the confidence of about 40-50% of phrase nodes in SST-3 task is above 0.9 and the accuracy of predicting these phrases is above 90% on the SST dataset. Considering the previous results, we speculate if we produce part of the phrase labels on generic texts, choose the predicted labels with high confidence and add them to the original SST training set during the training process, the results might be further improved.

6 Conclusion

We proposed SentiBERT, an architecture designed for capturing better compositional sentiment semantics. SentiBERT considers the necessity of contextual information and explicit syntactic guidelines for modeling semantic composition. Experiments show the effectiveness and transferability




Figure 6: The results of `SentiBERT` trained with part of the phrase-level labels on SST-5, SST-3 and Twitter Sentiment Analysis. We show the averaged results of 5 runs.

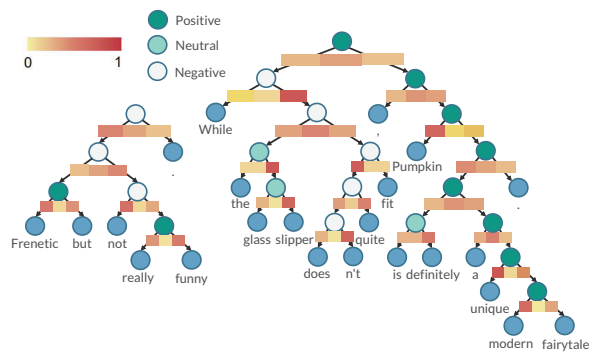


Figure 7: Cases for interpretability of compositional sentiment semantics. The three color blocks between parents and children are the attention weights distributed to left child, the phrase itself and right child.

of `SentiBERT`. Further analysis demonstrates its interpretability and potential with less supervision. For future work, we will extend `SentiBERT` to other applications involving phrase-level annotations.

Acknowledgement

We would like to thank the anonymous reviewers for the helpful discussions and suggestions. Also, we would thank Liunian Harold Li, Xiao Liu, Wasi Ahmad and all the members of UCLA NLP Lab for advice about experiments and writing. This material is based upon work supported in part by a gift grant from Taboola.

References

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450*.

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. In *LREC*, volume 10, pages 2200–2204.

Ankush Chatterjee, Kedhar Nath Narahari, Meghana Joshi, and Puneet Agrawal. 2019. Semeval-2019 Task 3: Emocontext Contextual Emotion Detection in Text. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 39–48.

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, Volume 1 (Long and Short Papers)*, pages 4171–4186.

Jacob Harer, Chris Reale, and Peter Chin. 2019. Tree-Transformer: A Transformer-Based Method for Correction of Tree-Structured Data. *arXiv preprint arXiv:1908.00449*.

Dan Hendrycks and Kevin Gimpel. 2017. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. *Proceedings of International Conference on Learning Representations*.

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. ACM.

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.

Pei Ke, Haozhe Ji, Siyang Liu, Xiaoyan Zhu, and Minlie Huang. 2019. SentiLR: Linguistic Knowledge Enhanced Language Representation for Sentiment Analysis. *arXiv preprint arXiv:1911.02493*.

Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014*

- Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 1746–1751. Association for Computational Linguistics.
- Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. 2017. Self-Normalizing Neural Networks. In *Advances in Neural Information Processing Systems*, pages 971–980.
- Bing Liu. 2012. Sentiment Analysis and Opinion Mining. *Synthesis Lectures on Human Language Technologies*, 5(1):1–167.
- Pengfei Liu, Xipeng Qiu, Xinchu Chen, Shiyu Wu, and Xuanjing Huang. 2015. Multi-timescale Long Short-Term Memory Neural Network for Modelling Sentences and Documents. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2326–2335.
- 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*.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Marco Marelli, Luisa Bentivogli, Marco Baroni, Raffaella Bernardi, Stefano Menini, and Roberto Zamparelli. 2014. Semeval-2014 Task 1: Evaluation of Compositional Distributional Semantic Models on Full Sentences through Semantic Relatedness and Textual Entailment. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 1–8.
- Jeff Mitchell and Mirella Lapata. 2008. Vector-based Models of Semantic Composition. In *Proceedings of ACL-08: HLT*, pages 236–244.
- Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. Semeval-2018 Task 1: Affect in Tweets. In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 1–17.
- Bo Pang and Lillian Lee. 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, page 271. Association for Computational Linguistics.
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs Up?: Sentiment Classification Using Machine Learning techniques. In *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing-Volume 10*, pages 79–86. Association for Computational Linguistics.
- Bo Pang, Lillian Lee, et al. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends® in Information Retrieval*, 2(1–2):1–135.
- Francis Jeffrey Pelletier. 1994. The Principle of Semantic Compositionality. *Topoi*, 13(1):11–24.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237.
- Fanchao Qi, Junjie Huang, Chenghao Yang, Zhiyuan Liu, Xiao Chen, Qun Liu, and Maosong Sun. 2019. Modeling Semantic Compositionality with Sememe knowledge. *ACL*.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. Semeval-2017 Task 4: Sentiment Analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 502–518.
- Richard Socher, Brody Huval, Christopher D Manning, and Andrew Y Ng. 2012. Semantic Compositionality through Recursive Matrix-Vector Spaces. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1201–1211. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1556–1566, Beijing, China. Association for Computational Linguistics.
- Yi Tay, Anh Tuan Luu, Siu Cheung Hui, and Jian Su. 2018. Attentive Gated Lexicon Reader with Contrastive Contextual Co-attention for Sentiment Classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3443–3453.

- Yaoshian Wang, Hung-Yi Lee, and Yun-Nung Chen. 2019. Tree Transformer: Integrating Tree Structures into Self-Attention. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1061–1070, Hong Kong, China. Association for Computational Linguistics.
- Hu Xu, Bing Liu, Lei Shu, and Philip S Yu. 2019. A Failure of Aspect Sentiment Classifiers and an Adaptive Re-weighting Solution. *arXiv preprint arXiv:1911.01460*.
- Xiaodan Zhu, Parinaz Sobhani, and Hongyu Guo. 2016. DAG-Structured Long Short-Term Memory for Semantic Compositionality. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 917–926, San Diego, California. Association for Computational Linguistics.

A Appendix

A.1 Details of Correlation Computation in Attention Networks

For vectors \mathbf{a} and \mathbf{b} , the correlation between them is computed as below:

$$\text{Attention}(\mathbf{a}, \mathbf{b}) = \tanh\left(\frac{1}{\alpha} \text{SeLU}((W_1 \times \mathbf{a})^T \times W_3 \times \text{SeLU}(W_2 \times \mathbf{b})),\right) \quad (3)$$

where SeLU (Klambauer et al., 2017) is an activation function and α equals 4. The two layers of attention networks do not share the parameters.

A.2 Details of Downstream Tasks

We adopt the following tasks for evaluation of sentence-level sentiment classifications:

SST-2,3 (Socher et al., 2013) These tasks all share with the text of the SST dataset and are single-sentence sentiment classification task, of which the numbers behind indicate the number of classes. Since two of five classes in SST-5 correspond to positive and another two indicate negative, with additional neutral ones, the dataset is separated into three groups in SST-3 task. We convert the 5-class phrase-level labels in SST-5 into three classes and leverage them in the training of SST-3 task.

Twitter Sentiment Analysis (Rosenthal et al., 2017) For Twitter Sentiment Analysis, given a tweet, model needs to decide which sentiment it expresses: positive, negative or neutral.

Emotion Intensity Ordinal Classification (Mohammad et al., 2018) The task is, given a tweet and an emotion, categorizing the tweet into one of four classes of intensity that best represents tweeter’s mental state. For Emotion Intensity Classification task, the metric is averaged Pearson Correlation value of the four subtasks, ‘happiness’, ‘sadness’, ‘anger’ and ‘fearness’.

Emotions in Textual Conversations (Chatterjee et al., 2019) In a dialogue, given a sentence with two turns of conversation, the models needs to classify the emotion expressed in the last sentence. For EmoContext, we follow the standard metrics used in Chatterjee et al. (2019) and use F1 score on the three classes ‘happy’, ‘sad’ and ‘angry’, except ‘others’ class, as the evaluation metric.

The statistics of datasets is shown in Table 5.

Dataset	Data Split	# of Classes
SST-phrase	8379 / 2184	5
SST-2	66475 / 859	2
SST-3	8379 / 2184	3
SST-5	8379 / 2184	5
Twitter	50284 / 12273	3
EmoContext	30141 / 2754	3
	sad: 1533 / 975	
EmoInt	angry: 1701 / 1001	4
	fear: 2252 / 986	
	joy: 1616 / 1105	

Table 5: Statistics of benchmarks.

Local Difficulty	0	1	2
Number	28136	10174	1342

Table 6: The distribution of nodes in terms of local difficulty.

Global Difficulty	0-4	5-9	10-14	15-19	20-23
Number	930	861	326	59	8

Table 7: The distribution of nodes in terms of global difficulty.

# of Negation Words	0	1	2-
Number	1825	325	34

Table 8: The distribution of nodes in terms of negation words.

Models	SST-phrase	SST-5
SentiBERT w/ token	68.23	56.02
SentiBERT w/ token and RoBERTa	68.78	56.91
SentiBERT	68.31	56.10
SentiBERT w/ RoBERTa	68.98	56.87

Table 9: The results after incorporating token node prediction. ‘Token’ denotes token node prediction.

A.3 Details of Analysis Part

The distribution of nodes and sentences in terms of local difficulty, global difficulty and negation words is shown in Table 6, 7 and 8, respectively.

A.4 Incorporating Token Node Prediction

Since the SST dataset also provides token-level sentiment labels, we combine the token node prediction with phrase node prediction learning objective together to model compositional sentiment semantics.

Results are shown in Table 9. We observe that the results drops a bit after additionally incorporating token-level sentiment information. This may be because the phrase sentiment is composed but the token sentiment mainly depends on the meaning of the lexicon itself rather than a kind of compositional sentiment semantics. The inconsistency of the training objectives may result in the performance drop.