

Investigating Typed Syntactic Dependencies for Targeted Sentiment Classification Using Graph Attention Neural Network

Xuefeng Bai ¹, Pengbo Liu ², and Yue Zhang ¹, *Member, IEEE*

Abstract—Targeted sentiment classification predicts the sentiment polarity on given target mentions in input texts. Dominant methods employ neural networks for encoding the input sentence and extracting relations between target mentions and their contexts. Recently, graph neural network has been investigated for integrating dependency syntax for the task, achieving the state-of-the-art results. However, existing methods do not consider dependency label information, which can be intuitively useful. To solve the problem, we investigate a novel relational graph attention network that integrates typed syntactic dependency information. Results on standard benchmarks show that our method can effectively leverage label information for improving targeted sentiment classification performances. Our final model significantly outperforms state-of-the-art syntax-based approaches.

Index Terms—Attention mechanism, dependency tree, graph neural networks, targeted sentiment analysis.

I. INTRODUCTION

TARGETED sentiment classification [1]–[5] is the task of predicting the sentiment polarity for target entity mentions in a given sentence. For example, suppose that a sentence is “I like the food here, but the service is terrible.” and the given targets are “food” and “service”. The output sentiment polarities on the two targets are *positive* and *negative*, respectively. Different from text level sentiment classification [6]–[8], targeted-sentiment is entity-centric, and therefore can offer more fine-grained opinion information from text documents. Dominant methods for targeted sentiment employ neural network models to encode the input sentence and target mention [9], [10]. Gates [3], [4], [11], convolution [12], [13], attention [14],

Manuscript received July 27, 2020; revised November 10, 2020; accepted November 10, 2020. Date of publication December 2, 2020; date of current version January 1, 2021. This work was supported by the National Natural Science Foundation of China under Grant 61976180. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Jianfeng Gao. (*Corresponding author: Yue Zhang.*)

Xuefeng Bai is with the Zhejiang University, Hangzhou 310007, China, and also with the School of Engineering, Westlake University, Hangzhou 310024, China (e-mail: xfbai.hk@gmail.com).

Pengbo Liu is with the Machine Intelligence and Translation Laboratory, School of Computer Science of Technology, Harbin Institute of Technology, Harbin 150001, China (e-mail: liupengbo.work@gmail.com).

Yue Zhang is with the School of Engineering, Westlake University, and also with the Institute of Advanced Technology, Westlake Institute for Advanced Study, Hangzhou 310024, China (e-mail: yue.zhang@wias.org.cn).

Digital Object Identifier 10.1109/TASLP.2020.3042009

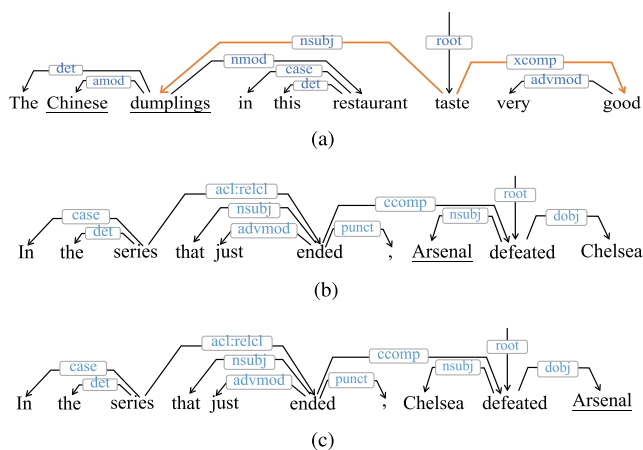


Fig. 1. (a) An example sentence with a dependency tree, (b, c) Two sentences which have similar dependency trees. The target mentions are underlined.

[15] and memory network [10], [16], [17] have been exploited to capture the relation between the target mention and its context information. The assumption is that deep syntactic and semantic features can be represented by neural encoding.

With the advance of structured neural encoders [2], [18]–[20], syntactic structures predicted by external parsers have shown their usefulness for the task. Intuitively, syntactic structures such as dependency trees can help better encode the correlation between a target mention and the relevant sentiment keywords. Recent methods consider dependency trees as adjacency matrices, using graph neural networks such as graph convolutional networks (GCN [19]) and graph attention network (GAT [20]) to encode the input sentence according to such matrices [21]–[23]. As shown in Fig. 1(a), such dependency tree structures can help bring target mentions closer to its relevant contexts, thereby facilitating feature representation. In this example, the relevant sentiment word “good” is distant from the target mention “Chinese dumplings” in the surface string, but close in the dependency tree (i.e., “Chinese dumplings” \xleftarrow{nsubj} “taste” \xrightarrow{xcomp} “good”).

While being more effective compared with the state-of-the-art approaches that do not use syntactic structures, these methods do not consider dependency labels, which can potentially be useful for sentiment disambiguation. As shown in Fig. 1(b) and

Fig. 1(c), while “Arsenal” has similar dependency arc relations with the word “defeated” in both sentence, the sentiment polarities are different. Apparently differentiating the label types can help differentiate such cases. On the other hand, considering arc label information can make the syntactic structure more sparse and thus increase the difficulty in learning. Thus it remains an open research question how to effectively make use of such fine-grained syntax features for better targeted sentiment classification.

We investigate a graph attention network to integrate such typed dependency features. In particular, the proposed model incorporates label features into the attention mechanism, using a novel extended attention function to guide information propagation from a target mention’s syntactic context to target mention itself. We name our model relational graph attention network (RGAT). With the help of these label features, our model can better capture the relationship between words, thus addressing the problems shown in Fig. 1. Moreover, the dependency labels can serve as additional features, enriching the representation of words.

Experiments over four benchmarks show that using GAT to encode the input gives better results compared to a Transformer encoder, which coincides with recent observation that syntax is useful for targeted sentiment classifications [21]–[23]. Further, adding typed dependency features gives consistently better results compared with a baseline without such information, thereby proving our research hypothesis. Our model gives significantly better results than state-of-the-art syntax-based work on the standard Laptop, Restaurant, Twitter and MAMS datasets. To our knowledge, we obtain the best reported results on all datasets without using external resources. Our code is released at <https://github.com/muyeb/RGAT-ABSA>.

II. RELATED WORK

With regard to structures, existing work on targeted sentiment classification can be classified into two main categories, namely those methods that do not rely on external syntax information and those using syntax structures.

A. Conventional Methods

Most work along the first line splits a given input sentence into two sections, including the target mention and its context. Features are extracted from each section and combining the features for making prediction. For example, Vo and Zhang [3] use word2vec embeddings and pooling mechanisms for extracting features from target mention and its left and right context, respectively, before concatenating all the features for classification. Zhang *et al.* [4] use gated recurrent neural network for extracting features from target mentions and its context, before further defining a gate to integrate such features. Tang *et al.* [11] exploit two long-short memory network (LSTM [24]) for feature encoding, and combine the last hidden state of two networks for classification.

There are also attempts exploiting convolutional neural networks (CNN) for targeted sentiment classification. For instance,

Huang and Carley [13] use parameterized filters and parameterized gates to incorporate aspect information into CNN, and apply the resulting CNN to encode the sentence. Xue and Li [25] further employ a gated CNN layer to extract aspect-specific features from the hidden states originated from a bi-directional RNN layer.

The attention mechanism [26], [27] is also shown to be useful for the task. Ma *et al.* [14] use a bidirectional attention network, which learns attention weights from the target mention to its context vectors, to model the target-context relationship. Li *et al.* [28] further improve the attention-based models with position information. Tang *et al.* [15] introduce a self-supervised attention learning approach, which automatically mines useful supervision information to refine attention mechanisms.

In addition to attention, Memory networks have also been applied to this task. Tang *et al.* [16] develop a deep memory network (MemNet), which uses pre-trained word vectors as a memory and exploits attention mechanism to update the memory. Chen *et al.* [17] improve MemNet by taking the hidden states generated by LSTM as memory and adopting gated recurrent units (GRU) to update the representation of target mentions. Wang *et al.* [10] deploy a targeted sensitive memory network for better information integration.

Recently, contextualized language models such as BERT [29], GPT [30] and ALBERT [31] ELECTRA [32] have shown their usefulness for a wide range of natural language inference (NLI) works. With regard to the task of targeted sentiment classification, Song *et al.* [33] use BERT to encode a target mention and its context, before applying attention to draw semantic interaction between targets and context words. Gao *et al.* [34] further introduce several variants to apply BERT for targeted sentiment classification, showing that incorporating target mention is helpful for BERT based models. Sun *et al.* [35] propose to construct an auxiliary sentence from the aspect and convert aspect-based sentiment classification into a sentence-pair classification task. Xu *et al.* [36] re-train a BERT model on a customer reviews dataset and use the resulting model for targeted sentiment classification. Li *et al.* [37] explore BERT for end-to-end targeted sentiment classification, jointly learning to predict the target mentions and its sentiment polarity.

It has been shown that BERT contains implicit syntactic information, which can be useful for downstream tasks [38], [39]. In contrast to these efforts, however, we make explicit use of syntactic information by encoding discrete structures. Our method is orthogonal to the implicit use of knowledge in BERT, and can be combined with contextualized embeddings.

B. Syntax-Based Methods

Among work that uses syntax structures, early work [40], [41] rely on manually-defined syntactic rules and use non-neural models. Subsequently, neural network models are explored for this task. For example, Dong *et al.* [2] use an adaptive recursive neural network (AdaRNN) to encode a syntax tree, transformed by placing the target mention onto the root node. Nguyen and Shirai [42] further extend AdaRNN into a phrase-level recursive neural network (Phrase AdaRNN), which takes both dependency

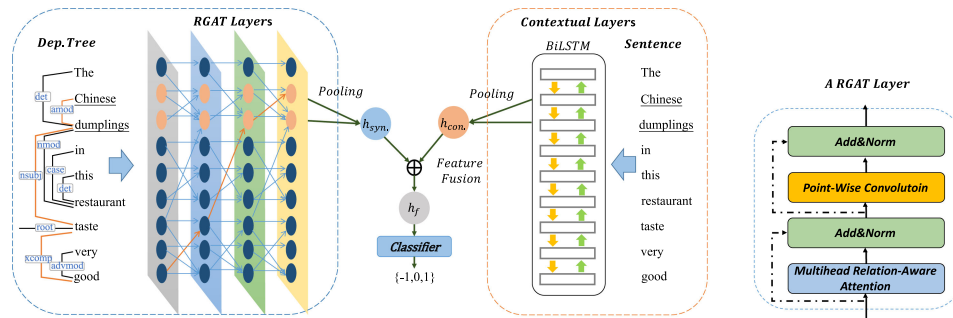


Fig. 2. Overview of the proposed network for targeted sentiment classification.¹ It consists of a contextual encoder (orange dotted frame), a RGAT encoder (blue dotted frame) and a classifier. The BiLSTM module can be replaced by BERT.

and constituent trees as input. He *et al.* [43] use the distance on a dependency tree to guide the attention mechanism, thus helping model to focus on more important context words.

Recent work uses graph neural networks to encode the syntactic structure, obtaining better results. In particular, Sun *et al.* [21] explore a GCN for encoding syntactical features to help the information exchange between target mention and context. Similar to that, Zhang *et al.* [23] introduce an aspect-specific GCN to incorporate the syntactical information and long-range word dependencies into the classification model. Huang and Carley [22] apply a GAT to model a dependency tree and integrate the GAT layer into LSTM to model the cross-layer dependency.

Our work is in this line. We take GAT as our base model and investigate a relational extension. Different from existing syntax-based work, we 1) exploit dependency relation information, which is proven to be useful by our work yet not investigated by [21]–[23]; 2) use an independent encoder for structure modeling instead of exploiting syntactic structures to refine contextual representations [21]–[23], [44]. Our double-encoder structure has two main advantages: a) the structure encoder is detachable and can be relatively easily applied to a new sequence encoder model (especially when the sentence length is not equal to the number of nodes in graph); b) it can reduce error propagation from automatic parsed dependency trees, as loss do not back-propagate from the tree representation to the sequence encoder.

Another similar attempt to this end is Shaw *et al.* [45], who extends a self-attention network (SAN) by integrating relative position information for neural machine translation. However, the relative position information is a simpler type of relation and thus is less informative. In contrast to their work, we investigate relational GAT for targeted sentiment classification, with significantly improved results.

III. APPROACH

For our model, each training instance consists of three components: a target mention, a sentence and a dependency tree of the sentence. Formally, we denote these components as a triplet: $\langle \mathcal{T}, \mathcal{S}, \mathcal{G} \rangle$, where $\mathcal{T} = \{w_i, w_{i+1}, \dots, w_{i+m-1}\}$ denotes a target mention word sequence, and $\mathcal{S} = \{w_1, w_2, \dots, w_i, \dots, w_{i+m}, \dots, w_n\}$ denotes a

sentence. The lengths of \mathcal{T} and \mathcal{S} are m and n , respectively. $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$ represents a syntactic graph (for example, a dependency tree) over \mathcal{S} , where \mathcal{V} includes all nodes (or words) $\{w_1, w_2, \dots, w_n\}$, \mathcal{A} is an adjacent matrix $\mathcal{A} \in \mathbb{R}^{n \times n}$ with $\mathcal{A}_{ij} = 1$ if there is a dependency arc relation between word w_i and w_j , and $\mathcal{A}_{ij} = 0$ otherwise, and \mathcal{R} is a label matrix, where \mathcal{R}_{ij} records the corresponding label of \mathcal{A}_{ij} if $\mathcal{A}_{ij} = 1$, and $\mathcal{R}_{ij} = \text{None}$ otherwise. The goal of targeted sentiment classification is to predict the sentiment polarity $y \in \{1, -1, 0\}$ of the sentence \mathcal{S} over the target mention \mathcal{T} , where 1, -1 and 0 denote *positive*, *negative* and *neutral*, respectively.

The overall architecture of our model is shown in Fig. 2. It is mainly composed of three components: the Contextual Encoder, the Relational Graph Attention (RGAT) Encoder and the Classifier. The Contextual Encoder can be viewed as a conventional model which applies Contextual encoder layers (e.g., BiLSTM, CNN, BERT) for feature learning and uses a simple pooling function for feature aggregation. The RGAT encoder is a syntax based encoder which incorporates syntax information into the progress of sentence modeling, thus generates syntax-aware word embeddings. The Feature Fusion mechanism is designed to dynamically combine the contextual and syntactic representations (denoted as h_{syn} and h_{con}). The final representation h_f is then fed into a Classifier for classification.

A. Contextual Encoder

We consider two types of models for contextual modeling: the first is a BiLSTM model, which is widely used for targeted sentiment classification and other tasks [26], [46]. The other is BERT [29], which is pre-trained on large-scale raw texts and has more parameters than BiLSTM.

BiLSTM We use BiLSTM to model the bidirectional context information. Following previous work [21], we use GloVe embeddings $v_i \in \mathbb{R}^{d_e}$, POS-tag embeddings $t_i \in \mathbb{R}^{d_t}$ as well as position embeddings $p_i \in \mathbb{R}^{d_p}$ as inputs features, where d_e , d_t and d_p denote the dimensions of word, POS-tag and position embeddings, respectively. Therefore, the representation of a word w_i is denoted as $x_i = [v_i; t_i; p_i]$, which is the concatenation of v_i , t_i and p_i .

¹In order to prevent the input graph from being too sparse, we remove the edge directions and add a self-loop edge for each word.

Given an word embedding sequence $x = \{x_1, x_2, \dots, x_n\}$, a forward $\overrightarrow{\text{LSTM}}$ generates a set of hidden states $\vec{h} = \{\vec{h}_1, \vec{h}_2, \dots, \vec{h}_n\}$, and a backward $\overleftarrow{\text{LSTM}}$ generates a set of hidden states $\overleftarrow{h} = \{\overleftarrow{h}_1, \overleftarrow{h}_2, \dots, \overleftarrow{h}_n\}$. Finally, the output hidden states $h = \{h_1, h_2, \dots, h_n\}$ are obtained by concatenating the corresponding forward and backward hidden states:

$$\begin{aligned} \vec{h} &= \overrightarrow{\text{LSTM}}([x_1, x_2, \dots, x_n]) \\ \overleftarrow{h} &= \overleftarrow{\text{LSTM}}([x_1, x_2, \dots, x_n]) \\ h &= [\vec{h}; \overleftarrow{h}] \end{aligned} \quad (1)$$

BERT is a pre-trained masked language model, which is based on a Transformer [27]. Previous work have shown that BERT can significantly boost the classification accuracy [33], [34], [37]. For comparability, we also adopt BERT to generate contextual embeddings.

To facilitate fine-tuning of the BERT model, we follow [33] to refactor the sequence as “[CLS]” + sentence + “[SEP]” + target mention + “[SEP]” as input to BERT. Formally, denoting the resulting sequence as²

$$x = \{x_0, x_1, \dots, x_n, x_{n+1}, x_{n+2}, \dots, x_{n+1+m}, x_{n+2+m}\}, \quad (2)$$

where x_0 and x_{n+1} are vector representations of “[CLS]” and “[SEP]” respectively. The BERT model generates a new sequence with the same length as x , denoted as:³

$$h = \{h_0, h_1, \dots, h_n, h_{n+1}, h_{n+2}, \dots, h_{n+1+m}, h_{n+2+m}\}, \quad (3)$$

where h_0 is called a “BERT pooling” vector, h_1, h_2, \dots, h_n are output contextual representations of the input word sequence. and $h_{n+2}, \dots, h_{n+1+m}$ corresponds to the output embedding of target mention. In this work, we use h_1, h_2, \dots, h_n for pooling and feature fusion.

Different from BiLSTM, the BERT based model does not take GloVe, POS-tag or position embeddings as additional input features, as such features are inherently learned by BERT itself.⁴

B. Relational Graph Attention Network

A relational graph attention network (RGAT) aims to perform information exchange among words according to the syntactic dependency paths. Compared with standard GAT networks, the RGAT method can additionally make use of the labeled relations (i.e., typed syntactic dependencies), thus generating more informative representations. In this section, we start by introducing the baseline GAT model which operates on an unlabeled graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, and then present RGAT, which extends the GAT with the ability to model a labeled graph $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{R})$.

1) *Vanilla Graph Attention Network*: We take the graph attention network as a baseline.⁵ The graph attention network [20]

²For simplicity, we assume that the BERT tokenizer [47] does not change the original tokenize results.

³Please refer to [29] for detailed BERT.

⁴We also tried to use such features, without observing significant improvements.

⁵We do not take the TD-GAT model of Huang and Carley [22] as our baseline, as TD-GAT contains a LSTM module which makes model more complicated

is a variant of graph neural networks, which leverages masked self-attention layers to encode graph structures. Compared with other types of graph neural networks (e.g. graph convolutional networks [19], graph recurrent networks [48]), GAT can be better interpreted thanks to the attention mechanism.

Input and output: A GAT takes a set of word (or node) embeddings $\{x_1, x_2, \dots, x_n\}$ as initial hidden states $\{h_1^0, h_2^0, \dots, h_n^0\}$, iteratively producing more abstract features $\{h_1^l, h_2^l, \dots, h_n^l\}$ with increasing $l, l \in [1, L]$. The l th GAT layer takes predecessor word features $\{h_1^{l-1}, h_2^{l-1}, \dots, h_n^{l-1}\}$ and an adjacent matrix \mathcal{A} as input, and produces a new set of word features $\{h_1^l, h_2^l, \dots, h_n^l\}$ as its output.

Feature aggregation: Given a word w_i with its neighbor word index $j \in \mathcal{N}(i)$,⁶ a GAT updates the word’s representation at layer l by calculating a weighted sum of the neighbor states. Briefly, the aggregation process of a multi-head-attention-based GAT can be described as:

$$h_i^l = \parallel_{z=1}^Z \sigma \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{lz} W_V^{lz} h_j^{l-1} \right), \quad (4)$$

where \parallel denotes vector concatenation, $W_V^{lz} \in \mathbb{R}^{\frac{d}{Z} \times d}$ is a parameter matrix of the z th head at layer l , d denotes the dimension of word feature vectors, Z is the number of attention heads, and σ represents the sigmoid activation function. The weight α_{ij}^{lz} models to what extent h_i^l depends on h_j^{l-1} :

$$\alpha_{ij}^{lz} = \frac{\exp(f(h_i^{l-1}, h_j^{l-1}))}{\sum_{j' \in \mathcal{N}(i)} \exp(f(h_i^{l-1}, h_{j'}^{l-1}))}. \quad (5)$$

If $j \notin \mathcal{N}(i)$, $\alpha_{ij}^{lz} = 0$. f is an attention function. We use the scaled dot-product attention function [27]:⁷

$$f(h_i^{l-1}, h_j^{l-1}) = \frac{(W_Q^{lz} h_i^{l-1})^T (W_K^{lz} h_j^{l-1})}{\sqrt{d/Z}}, \quad (6)$$

where $W_Q^{lz}, W_K^{lz} \in \mathbb{R}^{\frac{d}{Z} \times d}$ are parameter matrices of the z th head at layer l .

Apart from the aforementioned attention layer, we further add a point-wise convolution transformation (PCT) layer following the attention layer, which gives each node more information capacity. The convolution kernel size is 1, and convolution is applied to every single token belonging to the input. Given an output hidden states $h^l = \{h_1^l, h_2^l, \dots, h_n^l\}$ of the l th attention layer, the PCT layer is defined as:

$$PCT(h^l) = \delta(h^l \star W_{P_1} + b_{P_1}) \star W_{P_2} + b_{P_2}, \quad (7)$$

where δ refers to the ReLU activation function, \star denotes the convolution operation, $W_{P_1}, W_{P_2}, b_{P_1}, b_{P_2}$ are weights and bias of two convolutional kernels.

and harder to analyze. Besides, our preliminary experiments showed that our GAT baseline gives comparable results with TD-GAT.

⁶ $\mathcal{N}(i)$ is induced from the adjacent matrix \mathcal{A} .

⁷We also tried other attention functions, where no significant improvement founded.



Fig. 3. Attention distribution of RGAT model is a mixture of node-aware attention and relation-aware attention distribution. \otimes denotes the attention mixing operation (see Equation 10).

2) *Relational Graph Attention Network*: The vanilla GAT mentioned above uses an adjacent matrix as structural information, thus omitting dependency label features. RGAT incorporates relational features into attention calculation and aggregation process to obtain more informative representations. As shown in Fig. 3, the RGAT calculates two attention distributions, named as node-aware attention and relation-aware attention, taking their combination as final attention weights for feature aggregation.

Relations as input: Denoting the relation between word w_i and w_j as \mathcal{R}_{ij} , we transform \mathcal{R}_{ij} into a vector $r_{ij} \in \mathbb{R}^{d_r}$, where d_r is the dimension of relation embeddings. During the training process, the relation embeddings are jointly optimized with the model.

Relation-aware attention: Inspired by previous work on SAN [45], we consider relation features when calculating attentions weights between nodes. In practice, the RGAT method calculates an unnormalized node-aware attention e^N together with an unnormalized relation-aware attention e^R simultaneously. In particular, the node-aware attention e^N of the l th layer is the same as Equation (6):

$$e_{ij}^N = \begin{cases} f(h_i^{l-1}, h_j^{l-1}), & j \in \mathcal{N}(i) \\ -\text{inf}, & \text{otherwise} \end{cases} \quad (8)$$

while the relation-aware attention weight is given by:

$$e_{ij}^R = \begin{cases} f(h_i^{l-1}, r_{ij}), & j \in \mathcal{N}(i) \\ -\text{inf}, & \text{otherwise} \end{cases} \quad (9)$$

where r_{ij} denotes the vector representation of relation \mathcal{R}_{ij} . It should be noted that the relation embedding r_{ij} is shared among multiple layers and attention heads.

The above two types of attention scores are combined and normalized as:

$$\hat{\alpha}_{ij} = \frac{\exp(e_{ij}^N + e_{ij}^R)}{\sum_{j' \in \mathcal{N}(i)} \exp(e_{ij'}^N + e_{ij'}^R)}. \quad (10)$$

In this way, the resulting attention scores take both node features and relation features into consideration.

Relation-aware feature aggregation: Relations can also be important in the feature aggregation process, as these additional fine-grained information can be used to enrich the feature of each word h_i^l . To this end, we use the hidden feature vector of neighbor words h_j^{l-1} together with their corresponding relation

vector r_{ij} as inputs to update the representation of h_i^{l-1} :

$$h_i^l = \prod_{k=1}^Z \sigma \left(\sum_{j \in \mathcal{N}(i)} \hat{\alpha}_{ij}^{l,z} (W_V^{l,z} h_j^{l-1} + W_{V_r}^l r_{ij}) \right), \quad (11)$$

where $W_{V_r}^l \in \mathbb{R}^{\frac{d}{Z} \times d_r}$ is a parameter matrix.

In order to learn deep features, we apply a stacked relational graph attention network with multiple layers.

C. Pooling and Feature Fusion

With a contextual encoder (Section III-A) and a RGAT encoder (Section III-B), we obtain the contextual and syntax-aware target mention vectors, denoted as $\{h_i, h_{i+1}, \dots, h_{i+m-1}\}$ and $\{\hat{h}_i, \hat{h}_{i+1}, \dots, \hat{h}_{i+m-1}\}$, respectively. We then apply a pooling function over these vectors to obtain two global representation h_{con} and h_{syn} :

$$\begin{aligned} h_{con} &= \text{pool}(h_i, h_{i+1}, \dots, h_{i+m-1}), \\ h_{syn} &= \text{pool}(\hat{h}_i, \hat{h}_{i+1}, \dots, \hat{h}_{i+m-1}). \end{aligned} \quad (12)$$

In our implementation, *pool* is an average pooling function.⁸

In order to learn a composite representation which contains both contextual and syntax features, a fine-grained feature fusion mechanism is introduced to control the fusion ratio. We implement the feature fusion mechanism based on the gating mechanism [24], [49]. Formally, the syntax representation h_{syn} is fused into the contextual representation h_{con} by:

$$h_f = g \circ h_{syn} + (1 - g) \circ h_{con}, \quad (13)$$

where \circ is element-wise product operation, and g is a gate computed by:

$$g = \sigma(W_g[h_{syn}; h_{con}] + b_g). \quad (14)$$

In above equation, $[h_{syn}; h_{con}]$ is the concatenation of h_{syn} and h_{con} , W_g and b_g are model parameters.

D. Classifier

The classifier is a fully connected network, which takes the fused representation h_f as input and computes the probability of each sentiment class c :

$$P(y = c) = \frac{\exp(W h_f + b)_c}{\sum_{c' \in \mathcal{C}} \exp(W h_f + b)_{c'}}, \quad (15)$$

where W and b are tunable model parameters, and \mathcal{C} is the set of sentiment classes.

Given a set of training instances $D = \{d_1, d_2, \dots, d_N\}$, the training objective is a cross-entropy loss with L_2 regularization:

$$\ell = - \sum_{i=1}^N \sum_{c \in \mathcal{C}} I(y = c) \log(P(y = c)) + \lambda \|\Theta\|^2, \quad (16)$$

where I is an indicator function, N is the number of training examples, λ is a regularization hyperparameter and Θ denotes all the set of parameters in the model.

⁸We also tried other functions such as max pooling and random pooling, without observing significant improvements.

TABLE I
STATISTICS OF DATASETS

Dataset	Positive		Negative		Neutral	
	Train	Test	Train	Test	Train	Test
Restaurant	2164	727	807	196	637	196
Laptop	976	337	851	128	455	167
Twitter	1507	172	1528	169	3016	336
MAMS	3380	400	2764	329	5042	607

TABLE II
INPUT/OUTPUT DIMENSION OF MODEL SUB-MODULES

Name	GloVe	BERT
Embedding layer	1/300	1/768
Contextual encoder layer	360/100	768/768
RGAT encoder layer	100/100	100/100
Feature fusion layer	200/50	868/868
Classifier layer	50/3	868/3

IV. EXPERIMENTS

We conduct experiments on 4 benchmark datasets, including the Restaurant reviews (*Restaurant*) and Laptop reviews (*Laptop*) datasets of SemEval 2014 [50], the ACL14 *Twitter* dataset [2] and the MAMS dataset [51]. These datasets are labeled with three sentiment polarities: *positive*, *neutral* and *negative*. The number of samples in each category are summarized in Table I.

A. Settings

We adopt similar experimental settings as previous work [21], [22]. Two types of contextual encoders are considered: 1) The BiLSTM-based encoder; 2) the BERT-based encoder. With regard to the BiLSTM-based encoder, 300-dimensional GloVe vectors as adopted for word representation. The values are fixed during the training process. The dimensions of POS-tag, position and relation embeddings are set as 30. The dropout rate on input word embeddings is 0.7, and the L_2 regularization term $\lambda = 10^{-5}$. The Adamax [52] optimizer with a learning rate of 10^{-3} is adopted to train our models. For the BERT-based encoder, we adopt a pre-trained BERT⁹ for fine-tuning. The dropout rate on BERT embeddings is 0.1, and regularization term $\lambda = 10^{-5}$. The Adam [52] optimizer with a learning rate $2 * 10^{-5}$ is adopted for model training. For RGAT encoder, we use 5 attention heads in each layer. The input/output dimension of each layer is listed in Table II. We use the Deep Biaffine Parser [53]¹⁰ to obtain dependency trees.

We consider two metrics for model evaluation: Accuracy and Macro-Averaged F1. The latter is more appropriate for datasets with unbalanced classes. Pairwise t-test is conducted on both Accuracy and Macro-Averaged F1.

⁹We use pre-trained BERT-base-uncased model from [Online]. Available: <https://github.com/huggingface/transformers>

¹⁰[Online]. Available: <https://github.com/yzhangcs/parser>

B. Baselines

We compare our model with the state-of-the-art systems with-out and with syntactic knowledge. The syntax-free baselines include:

- **SVM** applies a support vector machine (SVM, [54]) model over extensive features for classification.
- **IAN** [55] represents the target and context interactively via two LSTMs and attention mechanism.
- **TNet** [12] transforms BiLSTM embeddings into target-specific embeddings, and uses a CNN for encoding.
- **MGAN** [56] exploits a BiLSTM to capture contextual information and a multi-grained attention mechanism to capture the relationship between aspect and context.
- **AOA** [57] introduces an Attention-over-Attention module to model the interaction between aspects and context sentences jointly.
- **AEN** [33] adopts attentional encoder network for feature representation and modeling semantic interactions between target and context.
- **CapsNet** [51] utilizes capsule networks [58] to model complicated relationships between target mentions and contexts.
- **BERT-PT** [36] uses a post-training approach on a pre-trained BERT model to improve the performance for review reading comprehension and targeted aspect sentiment classification.
- **BERT-SPC** [33] feeds “[CLS]” + sentence + “[SEP]” + target mention + “[SEP]” into a pre-trained BERT model, and then uses pooled embedding for classification.
- **CapsNet-BERT** [51] builds capsule networks on the top of BERT layers to predict sentiment polarities.

The syntax-based baselines are:

- **AdaRNN** [2] learns the sentence representation toward target via RNN semantic composition over a dependency tree.
- **PhraseRNN** [42] extends AdaRNN by adding two phrase composition functions. The PhraseRNN takes a dependency tree as well as a constituent tree as input.
- **SynAttn** [43] incorporates syntactic distance into the attention mechanism to model the interaction between target mention and context.
- **CDT** [21] and **ASGCN** [23] integrate dependency trees with GCN for aspect representation learning. Compared with CDT, ASGCN additionally apply attention mechanism to obtain final representation.
- **TD-GAT** [22] and **TD-GAT-BERT** [22] apply GAT to capture the syntax structure and improves it with LSTM to model relation across layers.

For fair comparison, we leave out baselines which rely on external resources such as auxiliary sentences [35] and other domain/language review corpus [34], [36], [59].

C. Main Results

Table III shows the results of different models. We first compare RGAT with the baseline (Transformer), which replace RGAT with a Transformer network. RGAT significantly

TABLE III

PERFORMANCE COMPARISON OF DIFFERENT MODELS ON THE BENCHMARK DATASETS. THE BEST PERFORMANCE ARE BOLD-TYPED, [†] DENOTES THAT THE MODEL REQUIRES DEPENDENCY TREE AS INPUT. RESULTS UNDERLINED INDICATE THAT THE PROPOSED METHOD IS SIGNIFICANTLY BETTER THAN STATE-OF-THE-ART MODEL AT SIGNIFICANCE LEVEL $p < 0.01$

Category	Model	Restaurant		Laptop		Twitter		MAMS	
		Accuracy	Macro-F1	Accuracy	Macro-F1	Accuracy	Macro-F1	Accuracy	Macro-F1
w/o Syn.	SVM	80.16	-	70.49	-	63.40	63.30	-	-
	IAN	78.60	-	72.10	-	-	-	76.60	-
	TNet	80.69	71.27	76.54	71.75	74.97	73.60	-	-
	MGAN	81.25	71.94	75.39	72.47	72.54	70.81	-	-
	AOA	81.20	-	74.5	-	-	-	77.26	-
	AEN	80.98	72.14	73.51	69.04	72.83	69.81	66.72	-
	CapsNet	80.79	-	-	-	-	-	79.78	-
	BERT-PT	84.95	76.96	78.07	75.08	-	-	-	-
	BERT-SPC	84.46	76.98	78.99	75.03	73.55	72.14	82.82	81.90
	AEN-BERT	83.12	73.76	79.93	76.31	74.71	73.13	-	-
	CapsNet-BERT	85.93	-	-	-	-	-	83.39	-
w Syn.	AdaRNN	-	-	-	-	66.30	65.90	-	-
	PhraseRNN	66.20	59.32	-	-	-	-	-	-
	SynAttn	80.45	71.26	72.57	69.13	-	-	-	-
	ASGCN	80.77	72.02	75.55	71.05	72.15	70.40	-	-
	CDT	82.30	74.02	77.19	72.99	74.66	73.66	80.70	79.79
	TD-GAT	81.20	-	74.00	-	-	-	-	-
	TD-GAT-BERT	83.00	-	80.10	-	-	-	-	-
Ours	Transformer	80.78	72.10	74.09	69.42	72.78	70.23	79.63	78.92
	RGAT	<u>83.55</u>	<u>75.99</u>	<u>78.02</u>	<u>74.00</u>	<u>75.36</u>	<u>74.15</u>	<u>81.75</u>	<u>80.87</u>
	Ours-RGAT-BERT	86.68	80.92	80.94	78.20	76.28	75.25	84.52	83.74

($p < 0.01$) improves the performance, with an accuracy improvement of 2.85 percent on average. Similarly, the average F1 score increases by 2.84 percent. This indicates that dependency trees can provide useful information for the targeted aspect sentiment classification task, which is consistent with observations by recent work [21], [22]. It is reasonable that Transformer gives weaker results than the state-of-the-art models, as the Transformer model adopts a simple pooling function rather than attention or CNN for target-context interaction modeling.

Compared with systems that do not rely on syntactic feature (SVM, IAN, TNet, MGAN, AOA, AEN, CapsNet), the RGAT model gives significantly ($p < 0.01$) better results. This observation is consistent with the results of our baseline. In addition, compared with recent work which takes dependency trees but without relation labels as input (AdaRNN, PhraseRNN, SynAttn, ASGCN, CDT, TD-GAT), RGAT also gives better results across all datasets. In particular, RGAT outperforms the state-of-the-art CDT model by 1.25 and 1.97 points on Restaurant with regard to accuracy and F1, respectively. Furthermore, the pre-trained BERT model can significantly boost the performance of each approach. With the help of BERT, the proposed model achieves better results than all the baselines, giving accuracy of 86.68, 80.94, 76.28 and 84.52 on Restaurant, Laptop, Twitter and MAMS, respectively. To our knowledge, we obtain the best reported results on all the datasets without using external resources.

V. ANALYSIS

A. Ablation Study

We present an ablation study on the effects of structure information and relation information. The experiments are conducted

TABLE IV

ABLATION STUDY ON THE RESTAURANT AND LAPTOP DATASETS. THE BEST PERFORMANCE ARE BOLD-TYPED. RATT REFERS TO RELATIONAL ATTENTION

Model	Restaurant		Twitter	
	Accuracy	Macro-F1	Accuracy	Macro-F1
Transformer	80.78	72.10	72.78	70.23
GAT	82.41	75.10	73.75	71.58
GAT-Ratt	82.68	75.39	74.60	73.05
RGAT	83.55	75.99	75.36	74.15
Transformer-BERT	84.89	77.90	73.43	72.08
GAT-BERT	85.70	78.80	74.67	73.50
GAT-Ratt-BERT	86.13	79.61	75.35	74.27
RGAT-BERT	86.68	80.92	76.28	75.25

on Restaurant, Twitter and MAMS as these three datasets have more training instances than Laptops. We consider three ablation baselines for comparison, including 1) Transformer: we replace the relational graph attention layers with self-attention layers. This model can be viewed as our model without dependency tree as supervision; and 2) GAT: the relation graph attention layers are substituted by graph attention layers, serving as our model without relation features; 3) GAT-Ratt: we remove the relation-aware feature aggregation module of RGAT to test the effectiveness of the relation-aware attention independently.

Effectiveness of structural information: As shown in Table IV, the GAT model gives consistently better accuracies and F1 scores over the Transformer model, with 1.3 and 1.2 percent average improvement of accuracy and F1 score, respectively. In addition, the performance of GAT-BERT is also better than that of Transformer-BERT. Such results indicates that explicit syntax knowledge is helpful for targeted sentiment classification. A possible reason for why improvement over BERT is relatively

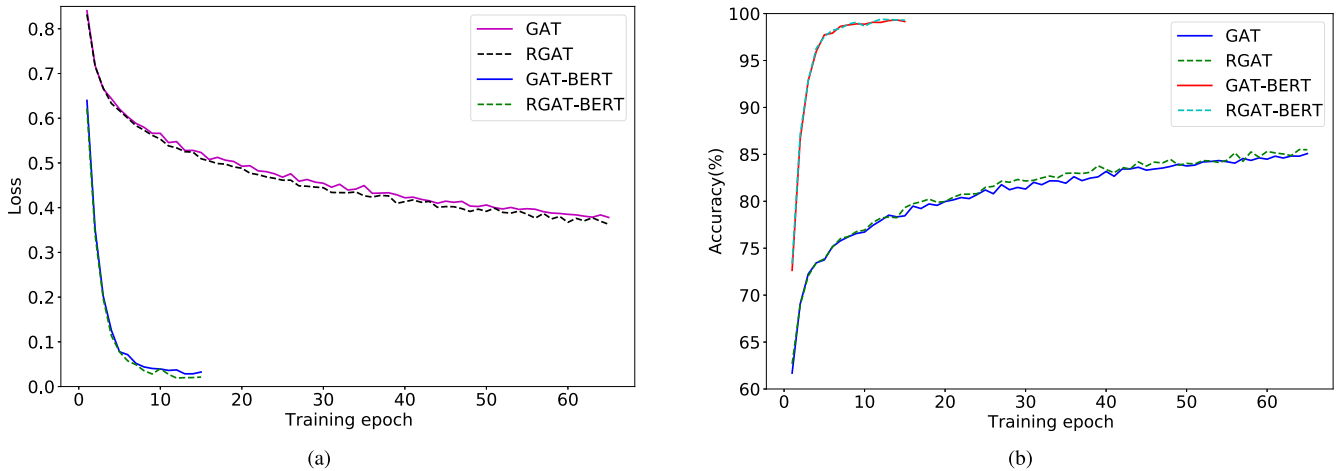


Fig. 4. Learning curve of GAT, RGAT, GAT-BERT and RGAT-BERT on MAMS.

small is that structural information is contained in the BERT representations to some extent because of contextual language modeling [60].

Effectiveness of relation-aware attention: GAT-Ratt consistently shows better performance over GAT. Specifically, GAT-Ratt outperforms GAT by 1.0 percent accuracy and 1.5 percent F1. Similar improvements are also observed for BERT based models. Such results indicate that the relation-aware attention mechanism is useful for targeted sentiment analysis.

Effectiveness of RGAT for typed dependencies: It can be observed that RGAT gives better results than GAT-Ratt and GAT on both datasets. In particular, RGAT outperforms GAT-Ratt on Twitter by 0.76 percent and 1.1 percent with regard to accuracy and F1, respectively. This observation confirms our intuition that the fine-grained relation information is helpful. Combined with relation-aware attention and relation-aware feature aggregation, the accuracy of GAT increases by 1.1 percent and 1.6 percent on Restaurant and Twitter, respectively. The improvement on F1 score reaches 0.9 percent and 2.6 percent, respectively. Similar trends can also be observed when comparing RGAT-BERT with GAT-BERT. Such results demonstrate that dependency labels can bring significant improvement to the baseline model, regardless of different input embeddings and contextual models.

We present the learning curves of GAT, RGAT, GAT-BERT and RGAT-BERT in Fig. 4. It can be seen that RGAT gives consistently higher accuracy and lower loss compared with GAT. Similarly, the results of RGAT-BERT are better than GAT-BERT. Such results prove that RGAT(-BERT) is more powerful than GAT(-BERT) from another perspective.

B. Impacts of Parsing Accuracy

We conduct experiments to study the effects of parsing accuracy on the classification performance. Specifically, we consider the following settings: 1) Transformer, which serves as a baseline model without dependency tree; 2) Random Tree, which uses a random tree as input for training;¹¹ 3) Random Label,

¹¹We select 10 random seeds and report the averaged results.

TABLE V
PERFORMANCE OF DIFFERENT PARSERS. UAS REFERS TO UNLABELED ATTACHMENT SCORE AND LAS REFERS TO LABELED ATTACHMENT SCORE

Model	UAS	LAS
Stanford Parser	94.10	91.49
Biaffine Parser	95.90	94.25
BERT Biaffine Parser	96.64	95.11

TABLE VI
THE ACCURACY AND F1 WHEN USING DIFFERENT DEPENDENCY TREES AS INPUT

Model	Restaurant		Twitter	
	Accuracy	Macro-F1	Accuracy	Macro-F1
Transformer	80.78	72.10	72.78	70.23
Random Tree	78.97	70.46	71.20	68.75
Random Labels	82.03	74.67	73.63	71.42
Stanford Parser	83.37	75.82	75.13	74.12
Biaffine Parser	83.55	75.99	75.36	74.15
BERT Biaffine Parser	83.64	75.91	75.28	74.10

which uses the gold tree structure (dependency heads) but randomly permutes the dependency labels; 4) Stanford Parser, for which dependency trees are parsed by Stanford Transition-based Parser [61]; 5) Biaffine Parser, which applies a Deep Biaffine Parser [53] to obtain dependency trees. 6) BERT Biaffine Parser, which adopts the BERT based Biaffine Parser to obtain a dependency tree. The performances of the Stanford Parser and Deep Biaffine Parser on the Penn Treebank are given in Table V.

Table VI gives the results of different systems on Restaurant and Twitter. It can be seen that a random tree has negative influences on the classification performance. The Random Label model gives better results than Transformer, showing that the tree structure is still useful for targeted sentiment classification, despite label noise. In addition, the trees parsed by pre-trained models (Stanford, Biaffine and BERT Biaffine parser) have more positive impact on model performance than random trees. This is likely because pre-trained parsers can give more accurate predictions and the parsed results are more consistent.

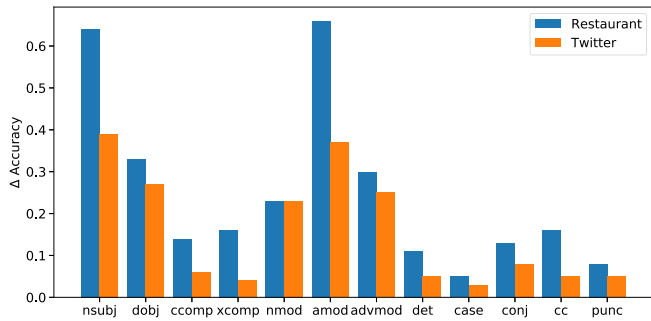


Fig. 5. Decrement of accuracy (compared with the baseline) on Restaurant and Twitter when removing different dependency labels.

Comparing the results of different parsers, it can be seen that both Biaffine Parser and BERT Biaffine Parser give better results than the Stanford Parser, and the performance of the Biaffine Parser and the BERT Biaffine Parser are comparable. Such results indicate that the classification performance is positively correlated with the parsing accuracy.

C. Effects of Different Dependency Labels

We select the most frequent dependency labels and study their contribution to the model performance. In particular, the following categories of dependency labels are considered: 1) Clausal Argument Relations, including $\{nsubj, dobj, ccomp, xcomp\}$; 2) Nominal Modifier Relations, including $\{nmod, amod, advmod, det, case\}$; 3) Other Notable Relations, including $\{conj, cc, punc\}$.

Fig. 5 gives the results of the RGAT model on the two benchmarks. We report the decrement of classification accuracy when removing different dependency labels independently. The accuracy of RGAT decreases most rapidly when removing *nsubj* and *amod* relations, indicating that *nsubj* and *amod* carry most important information for classification. This is intuitive as the *nsubj* label represents the nominal subject relation, which is especially informative when one side of this relation is a target mention. Similarly, the *amod* label denotes the adjective modifier relation, and is always related to the adjective which conveys the sentiment polarity. In addition, the *dobj*, *nmod* and *advmod* labels are also important factors to the model performance. The accuracy decreases by a range from 0.23 to 0.32 when these labels are removed. Last but not least, although having a high frequency, the *det*, *case* and *punc* dependency labels are the most irrelevant to the classification accuracy, with accuracy decline of about 0.07, 0.04 and 0.05, respectively. This indicates that the label frequency is not strongly correlated to the classification accuracy, and the relation label matters.

D. node Attention vs Edge Attention

In order to compare the importance of node attention and edge attention, we assign two trainable weighted factors to node attention and edge attention, respectively. In particular,




Fig. 6. Effects of model depth on classification accuracy on Restaurant.

we modified Equation 10 as:

$$\hat{\alpha}_{ij} = \frac{\exp(\beta_1 e_{ij}^N + \beta_2 e_{ij}^R)}{\sum_{j' \in \mathcal{N}(i)} \exp(\beta_1 e_{ij'}^N + \beta_2 e_{ij'}^R)}, \quad (17)$$

where $\beta_1, \beta_2 \in \mathbb{R}^1$ are model parameters.

We conduct evaluation on the Restaurant and Twitter datasets, and results are shown in Table VII. RGAT-weighted factors achieves comparable performance to RGAT, and RGAT-BERT-weighted factors gives slightly lower results compared with RGAT-BERT.

The values of β_1 and β_2 are shown in Table VIII. Interestingly, both values are close to 1, and this phenomenon is more obvious when using the BERT model. Such results can explain why the performance of RGAT-weighted factors is comparable to RGAT, and further indicate that node attention is of equal importance to relation attention.

E. Impacts of Model Depth

Fig. 6 shows the accuracy curves for RGAT and RGAT-BERT on Restaurant. Different numbers of layers ranging from 1 to 8 are considered. For RGAT, the initial accuracy is low and then increases along with the depth, reaching the best score of 83.55 with 6 layers. This is intuitive as the target-related sentiment words can be many-hops away from the target mention, which means that multiple layers of node communication is necessary for passing information from relevant context to the target mention. In contrast to RGAT, RGAT-BERT reaches the best accuracy (86.68) faster with 2 layers, which can be because that implicate syntax contained in BERT allows faster information propagation.

F. Case Study

We provide 4 examples to further analyze the proposed model. The attention scores over dependency edges are given for comparison.¹² Consider the case shown in Fig. 7(a) and Fig. 7(b). GAT gives a right sentiment in case 1 and a incorrect sentiment in case 2, while RGAT obtains the correct sentiment in both cases, which indicates that RGAT can differentiate similar structures, thanks to the use of dependency labels. In Fig. 7(c), GAT predicts

¹²For brevity, we omit the attention weight on self-loop edges.

TABLE VII
MODEL PERFORMANCE ON BENCHMARK DATASETS

Model	Restaurant		Twitter		MAMS	
	Accuracy	Macro-F1	Accuracy	Macro-F1	Accuracy	Macro-F1
Transformer	80.78	72.10	72.78	70.23	79.63	78.92
RGAT-weighted factors	83.01	76.35	75.51	73.84	81.52	80.68
RGAT	83.55	75.99	75.36	74.15	81.75	80.87
Transformer-BERT	84.89	77.90	73.43	72.08	82.23	81.47
RGAT-BERT-weighted factors	86.23	79.83	75.84	75.13	84.30	83.50
RGAT-BERT	86.68	80.92	76.28	75.25	84.52	83.74

TABLE VIII
THE VALUES OF TRAINABLE PARAMETER β_1 AND β_2 . WE REPORT MEAN \pm std VALUE OF MULTIPLE RGAT LAYERS

Dataset	RGAT		RGAT-BERT	
	β_1	β_2	β_1	β_2
Restaurant	0.992 \pm 0.19	1.194 \pm 0.17	1.000 \pm 5e $^{-5}$	1.000 \pm 3e $^{-5}$
Twitter	0.973 \pm 0.04	1.080 \pm 0.05	1.000 \pm 2e $^{-4}$	1.000 \pm 2e $^{-4}$
MAMS	0.925 \pm 0.15	1.093 \pm 0.22	0.993 \pm 5e $^{-4}$	1.000 \pm 3e $^{-4}$

TABLE IX
COMPARISON OF MODEL PARAMETERS, 1 M= 1e 6

Dataset	GloVe		BERT	
	BiLSTM	Ours	BERT-SPC	Ours
Restaurant	1.53M	1.90M	109.59M	110.00M
Laptop	1.25M	1.49M	109.58M	110.00M
Twitter	4.13M	4.50M	110.37M	113.97M
MAMS	2.71M	2.96M	109.59M	110.00M

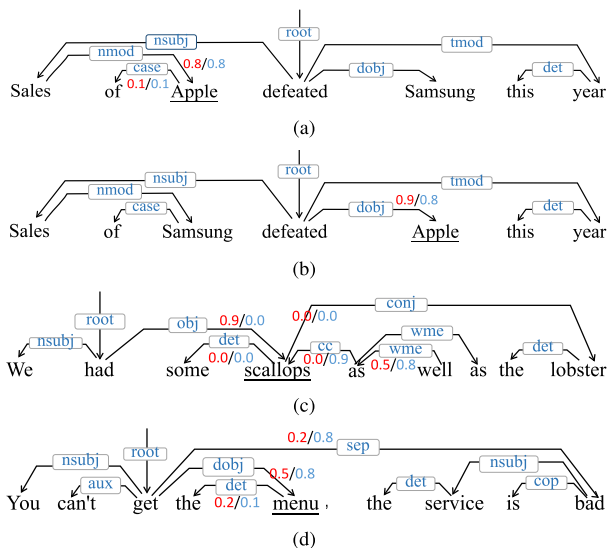


Fig. 7. Several samples, the numbers denote the attention weights given by RGAT (red) and GAT (blue). The target mentions are underlined.

a “positive” sentiment for target mention “scallops,” and the attention weight on the edge “as” \rightarrow “scallops” is high, which demonstrates that GAT is negatively influenced by “well,” which does not have positive meaning in this case. In contrast, RGAT gives the correct result and the attention weight on the edge “as” \xrightarrow{cc} “scallops” is close to zero. Fig. 7(d) shows another case where the decision of GAT is influenced by an irrelevant word “bad”. GAT’s attention weight on the edge “get” \xrightarrow{cc} “menu” is 0.8 while that of RGAT is 0.5. The last two examples indicate that the dependency labels have a positive influence on modeling the relationship between two words.

VI. DISCUSSION

Our approach is more complicated than previous syntax-free work, due to the modeling of typed dependencies. However, the

total number of parameters in RGAT does not increase significantly. The main reason is that we use a small hidden size (100) in RGAT encoder. Table IX shows the number of parameters of our method and two syntax-free systems (BiLSTM [11], BERT-SPC [33]).

In addition, our model requires more calculations than previous work. However, the encoding process of sentence and dependency tree are independent, so that the forward calculation and backward calculation of RGAT encoder and contextual encoder can be parallelized theoretically.

VII. CONCLUSION

We investigated the usage of typed dependency structures for targeted sentiment classification, by extending a graph attention network encoder with relation features. Extensive experiments on four standard benchmarks show that label information is useful for sentiment classification, and our relational GAT model can effectively encode such features. Our final model gives results that are better than the existing best results in the literature. We further study the impact of parsing performance, dependency labels and network depth on model performance, finding that different dependency arc labels do have different effects on targeted sentiment signal propagation, thereby motivating the use of arc label information. In addition, contextualized embeddings are complementary to structured dependencies for improving the results.

REFERENCES

- [1] L. Jiang, M. Yu, M. Zhou, X. Liu, and T. Zhao, “Target-dependent twitter sentiment classification,” in *Proc. ACL*, 2011, pp. 151–160.
- [2] L. Dong, F. Wei, C. Tan, D. Tang, M. Zhou, and K. Xu, “Adaptive recursive neural network for target-dependent twitter sentiment classification,” in *Proc. ACL*, 2014, pp. 49–54.
- [3] D.-T. Vo and Y. Zhang, “Target-dependent twitter sentiment classification with rich automatic features,” in *Proc. IJCAI*, 2015, pp. 1347–1353.
- [4] M. Zhang, Y. Zhang, and D.-T. Vo, “Gated neural networks for targeted sentiment analysis,” in *Proc. AAAI*, 2016, pp. 3087–3093.

- [5] S. Wang, S. Mazumder, B. Liu, M. Zhou, and Y. Chang, "Target-sensitive memory networks for aspect sentiment classification," in *Proc. ACL*, 2018, pp. 957–967.
- [6] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up? sentiment classification using machine learning techniques," in *Proc. EMNLP*, 2002, pp. 79–86.
- [7] A. Meena and T. V. Prabhakar, "Sentence level sentiment analysis in the presence of conjuncts using linguistic analysis," in *Proc. Adv. Inf. Retrieval*, G. Amati, C. Carpineto, and G. Romano, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 573–580.
- [8] A. Yessenalina, Y. Yue, and C. Cardie, "Multi-level structured models for document-level sentiment classification," in *Proc. EMNLP*, 2010, pp. 1046–1056.
- [9] Y. Wang, M. Huang, X. Zhu, and L. Zhao, "Attention-based LSTM for aspect-level sentiment classification," in *Proc. EMNLP*, 2016, pp. 606–615.
- [10] S. Wang, S. Mazumder, B. Liu, M. Zhou, and Y. Chang, "Target-sensitive memory networks for aspect sentiment classification," in *Proc. ACL*, 2018, pp. 957–967.
- [11] D. Tang, B. Qin, X. Feng, and T. Liu, "Effective LSTMs for target-dependent sentiment classification," in *Proc. COLING*, 2016, pp. 3298–3307.
- [12] X. Li, L. Bing, W. Lam, and B. Shi, "Transformation networks for target-oriented sentiment classification," in *Proc. ACL*, 2018, pp. 946–956.
- [13] B. Huang and K. Carley, "Parameterized convolutional neural networks for aspect level sentiment classification," in *Proc. EMNLP*, 2018, pp. 1091–1096.
- [14] D. Ma, S. Li, X. Zhang, and H. Wang, "Interactive attention networks for aspect-level sentiment classification," in *Proc. IJCAI*, 2017, pp. 4068–4074.
- [15] J. Tang *et al.*, "Progressive self-supervised attention learning for aspect-level sentiment analysis," in *Proc. ACL*, 2019, pp. 557–566.
- [16] D. Tang, B. Qin, and T. Liu, "Aspect level sentiment classification with deep memory network," in *Proc. EMNLP*, 2016, pp. 214–224.
- [17] P. Chen, Z. Sun, L. Bing, and W. Yang, "Recurrent attention network on memory for aspect sentiment analysis," in *Proc. EMNLP*, 2017, pp. 452–461.
- [18] K. S. Tai, R. Socher, and C. D. Manning, "Improved semantic representations from tree-structured long short-term memory networks," in *Proc. ACL*, 2015, pp. 1556–1566.
- [19] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [20] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [21] K. Sun, R. Zhang, S. Mensah, Y. Mao, and X. Liu, "Aspect-level sentiment analysis via convolution over dependency tree," in *Proc. EMNLP-IJCNLP*, 2019, pp. 5678–5687.
- [22] B. Huang and K. Carley, "Syntax-aware aspect level sentiment classification with graph attention networks," in *Proc. EMNLP-IJCNLP*, 2019, pp. 5468–5476.
- [23] C. Zhang, Q. Li, and D. Song, "Aspect-based sentiment classification with aspect-specific graph convolutional networks," in *Proc. EMNLP-IJCNLP*, 2019, pp. 4567–4577.
- [24] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, pp. 1735–1780, 1997.
- [25] W. Xue and T. Li, "Aspect based sentiment analysis with gated convolutional networks," in *Proc. ACL*, 2018, pp. 2514–2523.
- [26] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proc. ICLR*, 2015.
- [27] A. Vaswani *et al.*, "Attention is all you need," in *Proc. NIPS*, 2017, pp. 5998–6008.
- [28] L. Li, Y. Liu, and A. Zhou, "Hierarchical attention based position-aware network for aspect-level sentiment analysis," in *Proc. Conll*, 2018, pp. 181–189.
- [29] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL*, 2019, pp. 4171–4186.
- [30] A. Radford, "Improving language understanding by generative pre-training," 2018.
- [31] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "Albert: A lite bert for self-supervised learning of language representations," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [32] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning, "Electra: Pre-training text encoders as discriminators rather than generators," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [33] Y. Song, J. Wang, T. Jiang, Z. Liu, and Y. Rao, "Attentional encoder network for targeted sentiment classification," *CoRR*, vol. abs/1902.09314, 2019.
- [34] Z. Gao, A. Feng, X. Song, and X. Wu, "Target-dependent sentiment classification with bert," *IEEE Access*, vol. 7, pp. 154 290–154 299, 2019.
- [35] C. Sun, L. Huang, and X. Qiu, "Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence," in *Proc. NAACL*, 2019, pp. 380–385.
- [36] H. Xu, B. Liu, L. Shu, and P. Yu, "BERT post-training for review reading comprehension and aspect-based sentiment analysis," in *Proc. NAACL*, 2019, pp. 2324–2335.
- [37] X. Li, L. Bing, W. Zhang, and W. Lam, "Exploiting BERT for end-to-end aspect-based sentiment analysis," in *Proc. 5th Workshop Noisy User-Generated Text*, 2019, pp. 34–41.
- [38] Y. Goldberg, "Assessing bert's syntactic abilities," *CoRR*, vol. abs/1901.05287, 2019, *ArXiv1901.05287*.
- [39] K. Clark, U. Khandelwal, O. Levy, and C. D. Manning, "What does bert look at? an analysis of bert's attention," in *Proc. BlackBoxNLP@ACL*, 2019, pp. 276–286.
- [40] G. Qiu, B. Liu, J. Bu, and C. Chen, "Opinion word expansion and target extraction through double propagation," *Comput. Linguistics*, vol. 37, pp. 9–27, 2011.
- [41] K. Liu, H. Xu, Y. Liu, and J. Zhao, "Opinion target extraction using partially-supervised word alignment model," in *Proc. IJCAI*, 2013, pp. 2134–2140.
- [42] T. H. Nguyen and K. Shirai, "PhraseRNN: Phrase recursive neural network for aspect-based sentiment analysis," in *Proc. EMNLP*, 2015, pp. 2509–2514.
- [43] R. He, W. S. Lee, H. T. Ng, and D. Dahlmeier, "Effective attention modeling for aspect-level sentiment classification," in *Proc. 27th Int. Conf. Comput. Linguistics*, 2018, pp. 1121–1131.
- [44] K. Wang, W. Shen, Y. Yang, X. Quan, and R. Wang, "Relational graph attention network for aspect-based sentiment analysis," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics.*, Jul. 2020, pp. 3229–3238.
- [45] P. Shaw, J. Uszkoreit, and A. Vaswani, "Self-attention with relative position representations," in *Proc. NAACL*, 2018, pp. 464–468.
- [46] M. Sundermeyer, R. Schlüter, and H. Ney, "LSTM neural networks for language modeling," in *Proc. INTERSPEECH*, 2012, pp. 194–197.
- [47] Y. Wu *et al.*, "Google's neural machine translation system: Bridging the gap between human and machine translation," *CoRR*, vol. abs/1609.08144, 2016.
- [48] L. Song, Y. Zhang, Z. Wang, and D. Gildea, "A graph-to-sequence model for AMR-to-text generation," in *Proc. ACL*, 2018, pp. 1616–1626.
- [49] K. Cho *et al.*, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *Proc. EMNLP*, 2014, pp. 1724–1734.
- [50] M. Pontiki, D. Galanis, J. Pavlopoulos, H. Papageorgiou, I. Androutsopoulos, and S. Manandhar, "SemEval-2014 task 4: Aspect based sentiment analysis," in *Proc. SemEval*, 2014, pp. 27–35.
- [51] Q. Jiang, L. Chen, R. Xu, X. Ao, and M. Yang, "A challenge dataset and effective models for aspect-based sentiment analysis," in *Proc. Conf. Empirical Methods Natural Lang. Process. 9th Int. Joint Conf. Natural Lang. Process.*, Nov. 2019, pp. 6280–6285.
- [52] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [53] T. Dozat and D. C. Manning, "Deep biaffine attention for neural dependency parsing," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [54] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, 1995, pp. 273–297.
- [55] D. Ma, S. Li, X. Zhang, and H. Wang, "Interactive attention networks for aspect-level sentiment classification," in *Proc. IJCAI*, 2017, pp. 4068–4074.
- [56] F. Fan, Y. Feng, and D. Zhao, "Multi-grained attention network for aspect-level sentiment classification," in *Proc. EMNLP*, 2018, pp. 3433–3442.
- [57] B. Huang, Y. Ou, and K. M. Carley, "Aspect level sentiment classification with attention-over-attention neural networks," in *Proc. Int. Conf. Social Comput., Behav.-Cultural Model. Prediction Behav. Representation Model. Simul.*, 2018, pp. 197–206.
- [58] S. Sabour, N. Frosst, and G. E. Hinton, "Dynamic routing between capsules," in *Proc. Adv. Neural Inf. Process. Syst.*, I. Guyon *et al.* Eds., vol. 30. Curran Associates, Inc., 2017, pp. 3856–3866.
- [59] H. Yang, B. Zeng, J. hao Yang, Y. Song, and R. Xu, "A multi-task learning model for chinese-oriented aspect polarity classification and aspect term extraction," *Neurocomputing*, vol. 419, pp. 344–356, 2021.

- [60] G. Jawahar, B. Sagot, and D. Seddah, "What does BERT learn about the structure of language?," in *Proc. ACL*, 2019, pp. 3651–3657.
- [61] D. Chen and C. Manning, "A fast and accurate dependency parser using neural networks," in *Proc. EMNLP*, 2014, pp. 740–750.



Xuefeng Bai received the B.S. and M.S. degrees in computer science from the Harbin Institute of Technology, China, in 2017 and 2019, respectively. Since September 2019, he has been working toward the Ph.D. degree with Zhejiang University, Hangzhou, China. His research interests include sentiment analysis, machine translation, and text generation.



Pengbo Liu received the B.S. degree in computer science, in 2018 from the Harbin Institute of Technology, Harbin, China, where he is currently working toward the M.S. degree. His research interests include sentiment analysis, text generation, and natural language processing.



Yue Zhang (Member, IEEE) is currently an Associate Professor with Westlake University. He has been working on statistical parsing, text synthesis, natural language synthesis, machine translation, information extraction, sentiment analysis, and stock market analysis. His research interests include natural language processing and computational finance. He was the recipient of the Best Paper Awards of IALP 2017 and COLING 2018. He is on the editorial board for the *Transactions of Association of Computational Linguistics* as an Action Editor, *ACM Transactions on Asian and Low Resource Language Information Processing* as an Associate Editor, and *IEEE TRANSACTIONS ON BIG DATA* as an Associate Editor, and as Area Chairs of COLING 2014/18, NAACL 2015/19, EMNLP 2015/17/19, and ACL 2017/18/19. He gave conference tutorials at NAACL 2010, ACL 2014, and EMNLP 2016/18.