

An Iterative Multi-Knowledge Transfer Network for Aspect-Based Sentiment Analysis

Yunlong Liang^{1,2*}, Fandong Meng², Jinchao Zhang², Yufeng Chen¹,
Jinan Xu^{1†} and Jie Zhou²

¹Beijing Key Lab of Traffic Data Analysis and Mining,
Beijing Jiaotong University, Beijing, China

²Pattern Recognition Center, WeChat AI, Tencent Inc, China

{yunlongliang, chenyf, jaxu}@bjtu.edu.cn

{fandongmeng, dayerzhang, withtomzhou}@tencent.com

Abstract

Aspect-based sentiment analysis (ABSA) mainly involves three subtasks: aspect term extraction, opinion term extraction, and aspect-level sentiment classification, which are typically handled in a separate or joint manner. However, previous approaches do not well exploit the interactive relations among three subtasks and do not pertinently leverage the easily available document-level labeled domain/sentiment knowledge, which restricts their performances. To address these issues, we propose a novel **Iterative Multi-Knowledge Transfer Network (IMKTN)** for end-to-end ABSA. For one thing, through the interactive correlations between the ABSA subtasks, our IMKTN transfers the task-specific knowledge from any two of the three subtasks to another one at the token level by utilizing a well-designed routing algorithm, that is, any two of the three subtasks will help the third one. For another, our IMKTN pertinently transfers the document-level knowledge, *i.e.*, domain-specific and sentiment-related knowledge, to the aspect-level subtasks to further enhance the corresponding performance. Experimental results on three benchmark datasets demonstrate the effectiveness and superiority of our approach.

1 Introduction

Aspect-based sentiment analysis (ABSA) has drawn increasing attention in the community, which includes three subtasks: aspect term extraction (AE), opinion term extraction (OE) and aspect-level sentiment classification (SC). The first two subtasks aim to extract the aspect term and the opinion term appearing in one sentence, respectively. The goal of the SC subtask is to detect

the sentiment polarity towards the extracted aspect term.

Most existing studies generally handle each task separately (Tang et al., 2016; Wang et al., 2016b; Hu et al., 2019b) or take OE as auxiliary task for AE or SC (Wang et al., 2017; Li et al., 2018b; He et al., 2019), where these separate approaches need to be pipelined or integrated together for practical use. Recently, some researches point out that joint methods can achieve promising performance than separate ones, where only two subtasks are coupled, such as ⟨AE, OE⟩ (Wang et al., 2017; Dai and Song, 2019) or ⟨AE, SC⟩ (Luo et al., 2019; Zhou et al., 2019; He et al., 2019; Liang et al., 2021). More recently, Chen and Qian (2020) focus on modeling the interactive relations, *i.e.*, bidirectional $AE \leftrightarrow OE$, unidirectional $AE \rightarrow SC$ and unidirectional $OE \rightarrow SC$ with a collaborative learning framework. To further enhance these subtasks, several researchers seek to the external accessible document-level corpora (containing domain-specific/sentiment-related knowledge¹) due to the limited aspect-level data (Dai and Song, 2019; Chen and Qian, 2019; He et al., 2018, 2019). As a better case, He et al. (2019) merge the document-level domain-specific and sentiment-related knowledge together to enhance the AE and SC subtasks, where the two kinds of knowledge are indiscriminate.² Despite their effectiveness, we argue that the above methods are insufficient to yield satisfactory results for end-to-end ABSA task due to 1) they merely couple two subtasks or not modeling all bidirectional interactive relations among three subtasks ($AE \leftrightarrow OE$, $AE \leftrightarrow SC$ and $OE \leftrightarrow SC$), and 2) the document-level domain-specific/sentiment-related knowledge is coarsely used, which is insufficient

¹The two terms mean domain-relevant/sentiment-relevant linguistic knowledge, which are defined in (He et al., 2019)

²We denote it a coarse way to use the domain-specific knowledge and sentiment-related knowledge together. By contrast, a fine-grained way is to separately and pertinently exploit the two kinds of knowledge to expert their advantages.

*Work was done when Yunlong Liang was interning at Pattern Recognition Center, WeChat AI, Tencent Inc, China.

†Jinan Xu is the corresponding author.

to exert their advantages.

First, the interactive relations among three aspect-level subtasks are mutually collaborative. For instance, in the sentence “*The fish is very delicious.*”, the opinion term “*delicious*” indicates that the sentiment polarity of the aspect term “*fish*” is *positive*, suggesting the strong interactive correlation among them. Conversely, given the aspect term “*fish*” and its sentiment polarity *positive*, the word “*delicious*” rather than other words (e.g., “*very*”) in the sentence will be easily extracted as an opinion term. Therefore, the bidirectional relations between three aspect-level subtasks are closely related and they can incrementally promote one another, as shown in the left part of Fig. 1.

Second, the document-level corpora, which contain domain-specific and sentiment-related knowledge, should be pertinently utilized for enhancing the three aspect-level subtasks of ABSA. In fact, most aspect and opinion terms own distinct domain-specific properties (Peng et al., 2018) while sentiment polarities (i.e., *positive*, *negative*, and *neutral*) are typically domain-invariant. For instance, the aspect term “*fish*” and the opinion term “*delicious*” reflect distinct domain-specific characteristics, indicating that they belong to *Restaurant* domain rather than *Laptop* domain. Conversely, the domain-specific properties can help distinguish these aspect and opinion terms from other domains or background words (e.g., “*very*”). Therefore, the domain-specific knowledge should be pertinently leveraged to help identify the aspect term and the opinion term rather than on judging sentiment polarity. Meanwhile, the sentiment-related knowledge should be targeted at benefiting the SC subtask rather than the AE and OE subtasks, as shown in the right part of Fig. 1.

Therefore, we propose an **Iterative Multi-Knowledge Transfer Network (IMKTN)** to fully exploit the interactive relations via transferring knowledge at both the token level and the document level for the ABSA task. Partially inspired by the superiority of capsule network in distinguishing different features by feature clustering (Sabour et al., 2017), we design a novel routing algorithm, which can mutually transfer task-specific knowledge among the three aspect-level subtasks, as illustrated in the left part of Fig. 1. Furthermore, IMKTN employs a more fine-grained way to pertinently transfer document-level knowledge to aspect-level subtasks, as shown in the right part




Figure 1: The interactive relations among three aspect-level subtasks (the left) and two document-level subtasks (the right), which are explicitly modeled through knowledge transferring. Three aspect-level subtasks are highly semantic correlated, and thus can incrementally facilitate one another through task-specific knowledge transfer. The domain-specific knowledge from domain classification is only transferred to aspect and opinion term extraction, and the sentiment-related knowledge from document-level sentiment classification is only for aspect-level sentiment classification.

of Fig. 1, where the knowledge from domain classification subtask only serves for the AE and OE subtasks while the knowledge from document-level sentiment classification subtask only helps the SC subtask. All multi-knowledge transfer processes are iteratively conducted for fully exploiting the knowledge in all tasks to enhance the ABSA task.

In summary, our contributions are three-fold:

- We propose an iterative multi-knowledge transfer network for the ABSA task, which can well exploit the interactive relations via transferring the task-specific knowledge from any two of the three aspect-level subtasks to the third one for mutual promotion using a well-designed routing algorithm.
- We propose a more fine-grained way to pertinently transfer the document-level knowledge to further enhance the aspect-level tasks.
- Our approach³ significantly outperforms the existing methods and achieves new state-of-the-art results on three benchmark datasets, namely SemEval14 (Restaurant14 and Laptop14) (Pontiki et al., 2014) and SemEval15 (Restaurant15) (Pontiki et al., 2015).

2 Task Definition

In this section, we formulate the aspect-level tasks and document-level tasks, where the document-level tasks are taken as auxiliary tasks for improving the aspect-level tasks.

³The code is publicly available at: <https://github.com/XL2248/IMKTN>

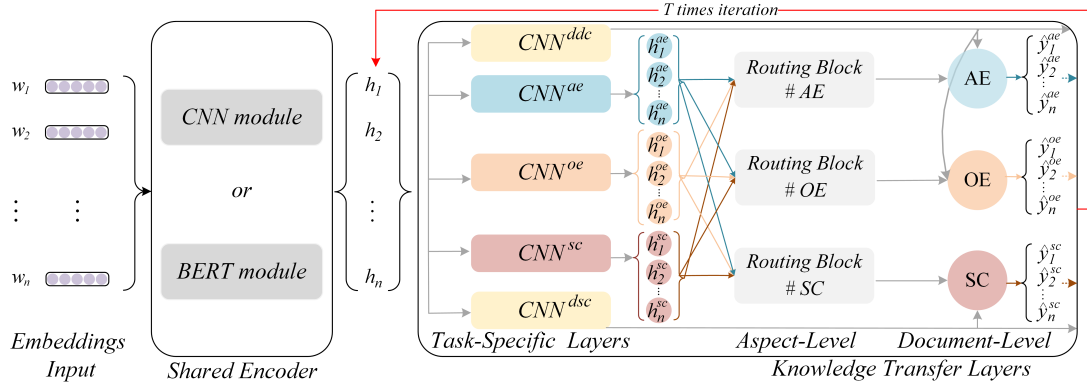


Figure 2: The model architecture of IMKTN. AE: aspect term extraction; OE: opinion term extraction; SC: aspect-level sentiment classification. To fully exploit the inter-task correlations among the three aspect-level subtasks for mutual promotion, the knowledge from them is mutually transferred to each other via the “Routing Block”. Besides, the knowledge from CNN^{ddc} is only transferred to the AE and OE subtasks. The knowledge from CNN^{dsc} is only transferred to the SC subtask. In summary, all the multi-knowledge transfer processes are iteratively conducted for adequately exploiting the knowledge from all the subtasks to enhance the ABSA task.

Aspect-Level Tasks. Following (Chen and Qian, 2020), the ABSA task is formulated as three sequence labeling subtasks. Given an input sentence $S = \{w_i\}_{i=1}^n$ with n words: 1) For the AE subtask, we aim to inference a tag sequence $Y^{ae} = \{y_i^{ae}\}_{i=1}^n$, where $y_i^{ae} \in \mathcal{Y}^{ae} = \{\text{BA}, \text{IA}, \text{O}\}$ denotes the **beginning** and the **inside** of an aspect term, and **other** words. 2) For the OE subtask, we aim to inference a tag sequence $Y^{oe} = \{y_i^{oe}\}_{i=1}^n$, where $y_i^{oe} \in \mathcal{Y}^{oe} = \{\text{BP}, \text{IP}, \text{O}\}$ denotes the **beginning** and the **inside** of an opinion term, and **other** words. 3) For the SC subtask, we aim to inference a tag sequence $Y^{sc} = \{y_i^{sc}\}_{i=1}^n$, where $y_i^{sc} \in \mathcal{Y}^{sc} = \{\text{pos}, \text{neg}, \text{neu}\}$ denotes *positive*, *negative* and *neutral* sentiment polarities.

Document-Level Tasks. This work contains two document-level subtasks: **domain classification** (DDC) and **sentiment classification** (DSC). For an input document $D = \{S_1, S_2, \dots, S_m\}$ with m sentences, the DDC and DSC aim to predict a domain label $Y^{ddc} \in \{\text{Laptop}, \text{Restaurant}\}$ and a sentiment label $Y^{dsc} \in \mathcal{Y}^{sc}$, respectively.

3 Model

As shown in Fig. 2, the IMKTN consists of four parts: 1) Shared Encoder, for extracting n-gram features; 2) Task-Specific Layers, for capturing sentence representations; 3) Aspect-Level Knowledge Transfer, including three Routing Blocks, for fully transferring knowledge among the aspect-level subtasks for mutual reinforcing; and 4) Document-Level Knowledge Transfer, for pertinently transfer-

ring document-level knowledge to corresponding aspect-level tasks. Finally, multi-source information is aggregated for the next iteration.

3.1 Shared Encoder

We apply two modules to extract sentence features, 1) we adopt Convolutional Neural Network (CNN) (Kim, 2014) as the feature extractor (Kalchbrenner et al., 2014); 2) we investigate a more powerful encoder (*i.e.*, BERT (Devlin et al., 2018)) as the backbone. The encoder is shared by the three aspect-level tasks and the two document-level tasks for providing common features.

3.2 Task-Specific Layers

Based on the Shared Encoder, 1) we design three aspect-level task-specific layers: CNN^{ae} , CNN^{oe} and CNN^{sc} , aiming to generate aspect-related knowledge, opinion-related knowledge, and sentiment-related knowledge, respectively; and 2) two document-level task-specific layers: CNN^{ddc} and CNN^{dsc} , for producing domain-specific features and sentiment features, respectively.

3.3 Aspect-Level Knowledge Transfer

As shown in Fig. 2, we design an aspect-level knowledge transfer layer, consisting of three Routing Blocks, to take full advantage of the inter-task knowledge among the three aspect-level subtasks.

Routing Block. The routing block serves for transferring knowledge among the aspect-level subtasks as shown in the “Routing Block” part of Fig. 2. Taking the “Routing Block #SC” for example, its

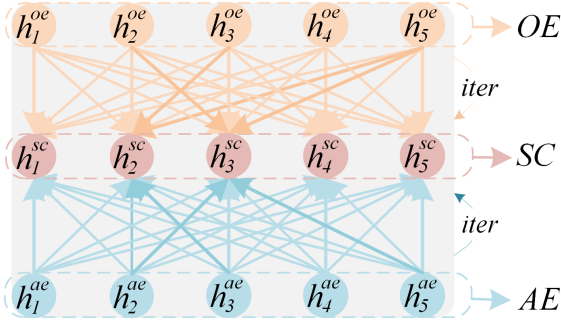


Figure 3: An example of the internal structure of “Routing Block #SC” as shown in Fig. 2. The knowledge of the AE and OE subtasks is transferred to the SC subtask through *iter* rounds of iteration, that is, the AE and OE subtasks will help the SC subtask.

internal structure is shown in Fig. 3, in which the knowledge from AE and OE is transferred to SC for enhancing its performance via our routing algorithm. We use the same algorithm to transfer knowledge from OE and SC to AE through the “Routing Block #AE”, from AE and SC to OE through the “Routing Block #OE”. In the conventional routing algorithm (Sabour et al., 2017), the high-level capsules are in a predefined fixed number, e.g., the total number of categories. While in our task, the high-level capsules are in dynamic numbers, where the number is determined by the sentence length. To this end, we propose a new routing algorithm, which is elaborated in detail below.

We show the whole routing process in Algorithm 1 by taking “transferring knowledge from OE to SC” as example. Specifically, the inputs of Algorithm 1 are the representation of OE ($\mathbf{h}_i^{oe} \in \mathbb{R}^{d_h}$) and iteration number (*iter*) (line 1). The $b_{j|i}$ is the probability indicating that the representation of the *i*-th token in OE agrees to be routed to the representation of the *j*-th token in SC, which is initialized with zero (line 2). The $\mathbf{W}^p \in \mathbb{R}^{n \times d_h \times d_o}$ is position-aware transformation matrix, which is realized via adding positional encoding (Vaswani et al., 2017), i.e., using $\text{AddPos}(\cdot)$ function to obtain the shared transformation matrix \mathbf{W} (line 3), where $\mathbf{W} \in \mathbb{R}^{d_h \times d_o}$. $PE_{(*)}$ is defined as:

$$PE_{(pos,2p)} = \sin(pos/10000^{2p/d_{model}}),$$

$$PE_{(pos,2p+1)} = \cos(pos/10000^{2p/d_{model}}),$$

where *pos* is token position in sentence, *p* is the positional index of the dimension and d_{model} is the input dimension. By doing so, the Algorithm can output capsules in dynamic numbers determined

Algorithm 1 Routing

```

1: procedure ROUTING ALGORITHM( $\mathbf{h}_i^{oe}, iter$ )
2:    $\forall i \in OE, \forall j \in SC, 1 \leq i, j \leq n, b_{j|i} \leftarrow 0.$ 
3:    $\mathbf{W}^p = \text{AddPos}(\text{tile}^a(\mathbf{W}, n), PE_{(pos,2p)}, PE_{(pos,2p+1)})$ 
4:    $\hat{\mathbf{u}}_{j|i} = \mathbf{h}_i^{oe} \mathbf{W}_{ij}^p$ 
5:   for iter iterations do
6:      $\forall i \in OE: \mathbf{c}_i \leftarrow \text{softmax}(\mathbf{b}_i)$ 
7:      $\forall j \in SC: \mathbf{s}_j \leftarrow \sum_i c_{j|i} \hat{\mathbf{u}}_{j|i}$ 
8:      $\forall j \in SC: \mathbf{v}_j^{oe} \leftarrow \text{squash}(\mathbf{s}_j)$ 
9:      $\forall i \in OE, \forall j \in SC: b_{j|i} \leftarrow b_{j|i} + \hat{\mathbf{u}}_{j|i} \cdot \mathbf{v}_j^{oe}$ 
10:  end for
11:  Return  $\mathbf{v}_j^{oe}$ 
12: end procedure

```

^aThe tile operation of Tensorflow (Abadi et al., 2016).

by the sentence length. The $\hat{\mathbf{u}}_{j|i}$ denotes the resulting opinion knowledge vector generated by multiplying the representation \mathbf{h}_i^{oe} with the specially-designed transformation matrix \mathbf{W}^p (line 4).

During each iteration (line 5), the coupling coefficients between low-level capsules \mathbf{h}_i^{oe} and high-level capsules \mathbf{v} are obtained by applying the softmax function (line 6). Then \mathbf{s}_j is calculated by aggregating all opinion vectors with $c_{j|i}$ as weights, voting for the sentiment polarity of the *j*-th token (line 7). After that the $\text{squash}(\mathbf{s}_j) = \frac{\|\mathbf{s}_j\|^2}{1+\|\mathbf{s}_j\|^2} \frac{\mathbf{s}_j}{\|\mathbf{s}_j\|}$ scales the output \mathbf{s}_j non-linearly to $0 \sim 1$ (line 8). Once the \mathbf{v}_j is updated in the current iteration, the probability $b_{j|i}$ becomes larger if the dot product $\hat{\mathbf{u}}_{j|i} \cdot \mathbf{v}_j^{oe}$ is large (line 9). That is, when the $\hat{\mathbf{u}}_{j|i}$ is more similar to the \mathbf{v}_j^{oe} , the dot product is larger, meaning that it is more likely to route this opinion knowledge to the *j*-th token and thus affects its sentiment polarity. Therefore, larger $b_{j|i}$ will lead to a larger agreement value $c_{j|i}$ between the opinion knowledge of the *i*-th token and the sentiment representation of the *j*-th token in the next iteration. In contrast, it generates low $c_{j|i}$ when there is no correlation between $\hat{\mathbf{u}}_{j|i}$ and \mathbf{v}_j^{oe} . After *iter* rounds of iteration, agreement values learned via the routing process ensure the opinion knowledge will be sent to the appropriate sentiment representation.

Similarly, we obtain the knowledge \mathbf{v}_j^{ae} , which is transferred from AE to SC, indicating which token should be correctly labeled with the sentiment polarity. Then the knowledge from AE and OE subtasks is combined as follows:

$$\mathbf{h}_j^{sc} = \text{Concat}(\mathbf{h}_j^{sc}, \mathbf{v}_j^{ae}, \mathbf{v}_j^{oe}),$$

where $\mathbf{h}_j^{sc} \in \mathbb{R}^{d_h+2d_o}$ is the *j*-th hidden state of the SC subtask (we set dimension size of all output capsules to d_o).

Through the process above, the multi-knowledge

transfer in “Routing Block #SC” is finished, which determines the sentiment polarity of each token in SC. Similarly, we achieve multi-knowledge transfer in “Routing Block #OE” and “Routing Block #AE” in Fig. 2. By doing so, three aspect-level subtasks are interacted with one another to fully exploit the inter-task correlations.

3.4 Document-Level Knowledge Transfer

We design the following two ways to pertinently transfer document-level knowledge to corresponding aspect-level tasks. (1) We transfer domain-specific knowledge ($a_i^{ddc(t)}$ and $a_i^{dsc(t)}$) from the DDC subtask to the AE and OE subtasks:

$$\mathbf{h}_i^{q(t+1)} = f_1([\mathbf{h}_i^{q(t)}; \hat{\mathbf{y}}_i^{ae(t)}; \hat{\mathbf{y}}_i^{oe(t)}; \hat{\mathbf{y}}_i^{sc(t)}; y_i^{ddc(t)}; a_i^{ddc(t)}]),$$

where $q \in \{ae, oe\}$, t is the iteration number ($0 \leq t \leq T$), $[\cdot; \cdot]$ denotes concatenation operation, $f_1(\cdot)$ is fully-connected layer and $\hat{\mathbf{y}}_i^{o(t)}$ is the prediction on the i -th token at the t -th iteration, which is proved helpful in (He et al., 2019), $o \in \{ae, oe, sc\}$. (2) We transfer sentiment-related knowledge ($\hat{\mathbf{y}}_i^{dsc(t)}$ and $a_i^{dsc(t)}$) from the DSC subtask to the SC subtask:

$$\mathbf{h}_i^{sc(t+1)} = f_2([\mathbf{h}_i^{sc(t)}; \hat{\mathbf{y}}_i^{ae(t)}; \hat{\mathbf{y}}_i^{oe(t)}; \hat{\mathbf{y}}_i^{sc(t)}; \hat{\mathbf{y}}_i^{dsc(t)}; a_i^{dsc(t)}]),$$

where $f_2(\cdot)$ is fully-connected layer. $a_i^{s(t)}$ ($s \in \{ddc, dsc\}$) is the self-attention weight (at the document level):

$$a_i^{s(t)} = \frac{\exp(\mathbf{h}_i^{s(t)} \mathbf{W}^s)}{\sum_{k=1}^n \exp(\mathbf{h}_k^{s(t)} \mathbf{W}^s)},$$

where \mathbf{W}^s is the trainable parameter. The document representation is computed by

$$\mathbf{h}^{s(t)} = \sum_{i=1}^n a_i^{s(t)} \mathbf{h}_i^{s(t)}.$$

Then a fully-connected layer with softmax function is applied to map $\mathbf{h}^{s(t)}$ to $\hat{\mathbf{y}}^{s(t)}$.

Overall, the IMKTN can fully perform knowledge transfer via the routing algorithm and pertinently incorporate the document-level knowledge to enhance the corresponding aspect-level tasks through such T rounds of iteration.

Datasets	Train			Test		
	#sent	#aspect	#opinion	#sent	#aspect	#opinion
D1 Restaurant14	3,044	3,699	3,484	800	1,134	1,008
D2 Laptop14	3,048	2,373	2,504	800	654	674
D3 Restaurant15	1,315	1,199	1,210	685	542	510

Table 1: Dataset statistics. #sent: sentences, #aspect: aspect terms and #opinion: opinion terms.

3.5 Training

For training, we minimize the loss on each token of aspect-level tasks and each instance of document-level tasks with the cross-entropy function. The aspect-level loss functions are written as follows:

$$\mathcal{J}_a = \lambda_1 \mathcal{L}_{ae} + \lambda_2 \mathcal{L}_{oe} + \lambda_3 \mathcal{L}_{sc},$$

$$\mathcal{L}_o = \frac{1}{n} \sum_{i=1}^n (\min(-\sum_{r=0}^{C_1} \mathbf{y}_{i,r}^o \log(\hat{\mathbf{y}}_{i,r}^{o(T)}))),$$

where λ_1, λ_2 and λ_3 are discount coefficients, $o \in \{ae, oe, sc\}$, n is the sentence length, C_1 is the class number, $\mathbf{y}_{i,r}^o$ denotes the ground-truth and $\hat{\mathbf{y}}_{i,r}^{o(T)}$ denotes the predictions with T times iteration. The document-level loss functions are formulated as follows:

$$\mathcal{J}_d = \lambda_4 \mathcal{L}_{ddc} + \lambda_5 \mathcal{L}_{dsc},$$

$$\mathcal{L}_s = \min(-\sum_{r=0}^{C_2} \mathbf{y}_r^s \log(\hat{\mathbf{y}}_r^{s(T)})),$$

where λ_4 and λ_5 are discount coefficients, $s \in \{ddc, dsc\}$, C_2 is the class number, \mathbf{y}_r^s denotes the ground-truth and $\hat{\mathbf{y}}_r^{s(T)}$ denotes the predictions after T times iteration.

For training the whole model, we firstly train the network with document-level tasks for a few epochs to generate reasonable features for aspect-level tasks. Then we train the network on the aspect-level and document-level corpus alternately, to minimize the corresponding loss.

4 Experiments

4.1 Experimental Settings

Datasets. We evaluate our model on three benchmark datasets from SemEval 2014 (Restaurant14 and Laptop14) (Pontiki et al., 2014) and SemEval 2015 (Restaurant15) (Pontiki et al., 2015), the data statistics of which is shown in Tab. 1. The opinion terms of these three datasets are annotated by Wang et al. (2016a). We adopt two document-level datasets from He et al. (2019), which include 30k instances of Yelp restaurant domain and 30k instances of Amazon electronic domain, respectively.

We merge the two datasets with domain labels for domain classification. We use the Yelp data when training on D1 and D3, and use the Amazon data for D2, due to the domain-specific properties.

Implementation Details. For fair comparison, we train our models with the same settings as comparison models (Chen and Qian, 2020). We tune the iteration number T and the routing number $iter$ on each validation set. More implementation and tuning details are given in Appendix A and B.

Evaluation Metrics. Following (Chen and Qian, 2020), four metrics are applied for evaluation, and the average score over 5 runs with random initialization is reported in all experiments. We use **F1-ae**, **F1-oe** and **F1-sc** to denote the F1-score of each subtask. We use F1-score denoted as **F1-absa** to measure the complete ABSA,⁴ where an extracted aspect term is taken as correct only when the span and the sentiment are both correct.

4.2 Comparison Models

To validate the performance of our proposed model on the ABSA task, we conduct contrast experiments with the following methods:

Pipeline Models. We respectively select two top performing models for AE: CMLA (Wang et al., 2017) and DECNN (Xu et al., 2018), and SC: TNet (Li et al., 2018a) and TCap (Chen and Qian, 2019), to construct 2×2 pipeline baselines. SPAN-BERT (Hu et al., 2019b) utilizes $BERT_{LARGE}$ as backbone networks for AE and SC subtasks.

Integrated Models. MNN (Wang et al., 2018) and INABSA (Li et al., 2019a): Both models handle the aspect term-polarity co-extraction as a sequence labeling problem with a unified tagging scheme.

Joint Models. The joint models including DOER (Luo et al., 2019), Span-based (Zhou et al., 2019), IMN (He et al., 2019), DREGCN (Liang et al., 2021), and RACL (Chen and Qian, 2020) are used to compare with ours, which are introduced in § 1 part.

For fair comparison, we validate IMKTN based on two encoders. 1) Based on CNN, we use GloVe embeddings (Pennington et al., 2014) and denote it as IMKTN-GloVe. 2) Based on

⁴Following (Chen and Qian, 2020), we use the predicted sentiment of the first word as the SC result if an aspect term has multiple words. Besides, aspect terms with *conflict* sentiment labels are ignored. All baseline models apply the same setting for fair comparison.

$BERT_{LARGE}$ (Devlin et al., 2018), we fine-tune it for ABSA, denoted as IMKTN-BERT).

4.3 Main Results

Results in Tab. 2 are divided into four groups: M1~M4, M5~M6, and M7~M12 are GloVe-based pipeline, integrated, and joint models, respectively. M13~M17 are BERT-based models.

1) Among all GloVe-based models (M1~M12), our IMKTN-GloVe significantly surpasses other baselines in most cases, and achieves 2.13%, 2.33%, and 2.86% absolute gains over RACL in terms of the overall metric F1-absa on three datasets. This suggests that the inter-task correlations and document-level knowledge have an overall positive impact on these subtasks, and demonstrates the superiority of our model. Furthermore, IMKTN-GloVe also obtains the best or the second best results on all subtasks, which further shows the effectiveness of our model. Another observation is that the joint models (M7~M12) perform better than pipeline and integrate models (M1~M6).

2) All BERT-based models get higher results than GloVe-based models thanks to the large-scale external knowledge (M13~M17 vs. M1~M12). Among all BERT-based Models, we observe that IMKTN-BERT significantly outperforms other BERT-based models, which suggests the effectiveness of our approach by transferring multi-source knowledge even based on the strong baseline and yields new state-of-the-art results on most metrics.

5 Analysis and Discussion

5.1 Whether Three Aspect-Level Subtasks Promote Each Other?

We evaluate the aspect-opinion pair F1 and aspect-opinion-sentiment triplet F1 on the test set (Fan et al., 2019; Peng et al., 2020; Xu et al., 2020b),⁵ for verifying whether the multi-knowledge transferring can help each other. The results are shown in Tab. 3, where IMKTN-D denotes removing all document-level knowledge transferring. We can see that our IMKTN-D can surpass the comparison models by a large margin under two settings. Particularly, in the aspect-opinion-sentiment triplet setting, IMKTN-D significantly outperforms other baselines, suggesting that inter-task knowledge transferring has an overall positive impact on these

⁵Table 9 (in Appendix) shows an example for explaining what are aspect-sentiment pair, aspect-opinion pair, and aspect-opinion-sentiment triplet.

Models	Restaurant14 (D1)				Laptop14 (D2)				Restaurant15 (D3)			
	F1-ae	F1-oe	F1-sc	F1-absa	F1-ae	F1-oe	F1-sc	F1-absa	F1-ae	F1-oe	F1-sc	F1-absa
M1 CMLA-TNet*	81.91	83.84	69.69	64.49	77.49	76.06	68.30	55.94	67.73	70.56	62.27	55.00
M2 CMLA-TCap*	82.45	82.67	72.23	65.34	76.80	77.33	69.52	55.56	68.55	71.07	66.45	55.47
M3 DECNN-TNet*	82.79	-	70.45	65.80	79.38	-	68.69	57.39	68.52	-	62.41	55.69
M4 DECNN-TCap*	82.79	-	71.77	66.84	79.38	-	69.61	57.71	68.52	-	63.60	56.22
M5 MNN*	83.05	84.55	68.45	63.87	76.94	77.77	65.98	53.80	70.24	69.38	57.90	56.57
M6 INABSA*	83.92	84.97	68.38	66.60	77.34	76.62	68.24	55.88	69.40	71.43	58.81	57.38
M7 DOER*	84.63	-	64.50	68.55	80.21	-	60.18	56.71	67.47	-	36.76	50.31
M8 Span-based	84.13	-	69.73	68.22	78.43	-	69.77	57.57	69.96	-	59.95	58.97
M9 IMN [‡]	83.33	85.61	75.66	69.54	77.96	77.51	72.02	58.37	70.04	71.94	71.76	59.18
M10 DREGCN [‡]	85.93	86.05	73.32	70.21	79.45	75.40	73.46	61.60	71.00	70.55	73.35	61.06
M11 RACL*	85.37	85.32	74.46	70.67	81.99	79.76	71.09	60.63	72.82	78.06	68.69	60.31
M12 IMKTN-GloVe	87.91[†]	87.65[†]	76.66[†]	72.80[†]	83.19[†]	81.82[†]	74.93[†]	62.96[†]	74.96[†]	74.48	75.39[†]	63.17[†]
M13 SPAN-BERT*	86.71	-	71.75	73.68	82.34	-	62.50	61.25	74.63	-	50.28	62.29
M14 IMN-BERT*	84.06	85.10	75.67	70.72	77.55	81.00	75.56	61.73	69.90	73.29	70.10	60.22
M15 DREGCN-BERT [‡]	87.00	86.95	75.79	72.60	79.78	79.21	76.37	63.04	73.30	72.60	73.02	62.37
M16 RACL-BERT*	86.38	87.18	81.61	75.42	81.79	79.72	73.91	63.40	73.99	76.00	74.91	66.05
M17 IMKTN-BERT	87.13[†]	88.62[†]	81.35	76.75[†]	83.89[†]	81.90[†]	76.42[†]	65.74[†]	74.63	76.79[†]	76.85[†]	68.33[†]

Table 2: Model comparison. We separate the results into the GloVe-based (M1~M12) and BERT-based (M13~M17) methods for fair comparison. Following RACL (Chen and Qian, 2020), we report average results over 5 runs with random initialization. The results with the symbol “*” refer to RACL. “[‡]” indicates that the results are referred to the original paper. “[†]” denotes our method is statistically significant (Koehn, 2004) better than RACL (p -value < 0.05), which is the best previous model.

	Models	D1	D2	D3
Aspect-Opinion Pair	IMN [‡]	54.94	54.87	56.45
	DREGCN [‡]	53.76	54.89	55.23
	RACL [‡]	54.67	54.75	56.74
	IMKTN-D	56.74[†]	56.60[†]	58.32[†]
Aspect-Opinion-Sentiment Triplet	IMN [‡]	50.95	41.21	45.65
	DREGCN [‡]	49.32	41.97	44.38
	RACL [‡]	50.65	41.55	45.45
	IMKTN-D	52.45[†]	44.82[†]	48.50[†]

Table 3: F1 scores (%). The aspect-sentiment pair results are shown in Tab. 2, *i.e.*, F1-absa score. “[‡]”: results are generated by running their official code. “[†]”: significantly better than RACL (p -value < 0.05).

#	Methods	F1-ae	F1-oe	F1-sc
0	Coarse way	81.06	85.02	65.44
1	Fine-Grained way	82.25	86.36	68.80

Table 4: F1 (%) on the validation set of D1.

aspect-level subtasks and hence the aspect-level subtasks indeed can promote each other.

5.2 Whether Pertinently Transferring Document-Level Knowledge Helps Aspect-Level Subtasks More?

In Tab. 4, the “Coarse way” (He et al., 2019) indicates that the knowledge from DDC and DSC is merged to indistinguishably enhance all aspect-level tasks. By contrast, the “Fine-Grained way” is to pertinently transfer the knowledge, *i.e.*, the

#	Models	D1	D2	D3
0	w/o AE KT	1.05/1.98↓	1.56/1.44↓	1.45/3.45↓
1	w/o OE KT	0.98/0.45↓	0.96/0.52↓	1.13/2.09↓
2	w/o SC KT	1.89/2.78↓	1.83/2.23↓	2.38/4.54↓
3	w/o DDC	1.88/2.03↓	1.87/1.82↓	1.54/3.37↓
4	w/o DSC	2.37/2.77↓	2.13/2.43↓	2.87/5.04↓

Table 5: Ablation study. “↓” denotes a performance drop of “IMKTN-GloVe/IMKTN-BERT” on the validation set (F1-absa). “KT”: knowledge transferring.

knowledge from DDC only transferred to AE and OE subtasks, and the knowledge from DSC only transferred to SC subtask. The results show that pertinently transferring document-level knowledge helps aspect-level subtasks more, which is consistent with our intuition that the domain-specific knowledge prefers to promote the AE and OE subtasks, and the sentiment-related knowledge tends to improve the SC subtask. Therefore, a fine-grained way is very necessary to enhance the ABSA.

5.3 Ablation Study

Tab. 5 shows the impact of different knowledge, where we remove one knowledge at a time. We conclude that: (1) once any of the aspect-level subtask knowledge transfer is removed (rows 0~2), scores on three benchmark datasets decrease under the both setting (*i.e.*, GloVe and $BERT_{LARGE}$), showing that the three aspect-level subtasks are

highly semantically correlated and thus can incrementally boost one another. (2) we also observe obvious drops when removing the document-level knowledge, especially when the DSC subtask is removed, suggesting that pertinently transferring the document-level knowledge significantly benefits the corresponding aspect-level tasks (rows 3~4).

5.4 Why using Capsule Network?

In our preliminary experiments, we conduct some experiments to investigate how to effectively transfer knowledge between different tasks. The results are shown in Tab. 6, where the capsule network (row 3) performs the best. The reason is capsules in adjacent layers connected by dynamic routing, which has the ability to distinguish different features by feature clustering (Sabour et al., 2017). This coincides with our motivation, *i.e.*, transferring related features from two subtasks to the third one through the bidirectional interactive relations for mutual promotion (feature clustering). However, other methods (rows 0~2) have no such dynamic routing mechanism and thus cannot dynamically conduct feature extraction and clustering, leading to unsatisfactory results. Therefore, we select the capsule network.

5.5 Case Study and Visualization

To provide an understanding of how the multi-knowledge transfer works, in Fig. 4,⁶ we take the knowledge transfer from OE and AE to SC for example to visualize the agreement value $c_{j|i}$. Fig. 4(a) and Fig. 4(c) are the cases of transferring knowledge from OE to SC. Fig. 4(a) shows that the knowledge of opinion term “longer” from the OE subtask is mainly sent to aspect term “battery” of the SC subtask and Fig. 4(c) shows the same phenomenon (the knowledge of opinion term “not terrible” from the OE subtask is mainly sent to the aspect term “prices”) though it is a negation sentence, indicating that the opinion word affects the sentiment polarity of the aspect term, *i.e.*, the former (AE) is naturally correlated with the latter (SC). Particularly, in Fig. 4(c), negation information can be effectively transferred to the aspect term “prices” via the routing algorithm and affects its sentiment polarity. Fig. 4(b) and Fig. 4(d) are the cases of transferring knowledge from AE to SC, showing that the aspect-related knowledge is mainly transferred to the aspect term “battery” and “prices”,

⁶Both two examples are taken from the Laptop 14 dataset.

#	Methods	D1	D2	D3
0	Concat	60.56	50.11	67.73
1	LSTM	60.77	51.19	66.93
2	Attention	61.36	52.49	68.02
3	Capsule	62.89	54.10	70.36

Table 6: F1-absa (%) on the validation set. Apart from using “Routing Blocks” in Fig. 2, we also try the following three methods. i): We directly concatenate the task-specific features (row 0). ii) We use an LSTM to sequentially read the task-specific features for transferring knowledge (row 1). iii) We apply attention to calculate the score between the task-specific features, and then take the score as the weight to conduct the task-specific knowledge transferring (row 2).

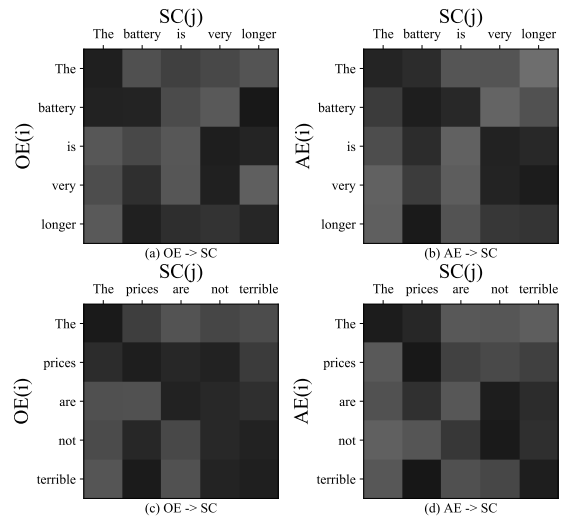


Figure 4: Visualization of $c_{j|i}$. The darker the color is, the more knowledge is transferred.

voting for them to be aspect terms. Therefore, the AE subtask can help the aspect-level sentiment classification to judge whether the word should own sentiment polarity or not. Besides, we also present thorough error analysis in Appendix C.

6 Related Work

Aspect-Based Sentiment Analysis. Existing models typically handle the ABSA task independently or jointly. Apparently, separately treating each subtask cannot exploit the inter-task correlations, leading to restricted performances, such as AE (Qiu et al., 2011; Liu et al., 2013, 2014, 2015; Yin et al., 2016; Li and Lam, 2017; Li et al., 2018b; Angelidis and Lapata, 2018; Ma et al., 2019, etc) and SC (Dong et al., 2014; Nguyen and Shirai, 2015; Vo and Zhang, 2015; Chen et al., 2017; Wang et al., 2018; Ma et al., 2018; Hu et al., 2019a; Liang et al., 2019; Bao et al., 2019; Sun et al., 2019; Tang

et al., 2019,?; Xu et al., 2020a, etc). By contrast, the integrated or joint methods (Wang et al., 2016a; Mitchell et al., 2013; Zhang et al., 2015; Li and Lu, 2017; Schmitt et al., 2018; Li et al., 2019b; Lin and Yang, 2020; Liang et al., 2021; Chen and Qian, 2020) can model the interactive correlations and thus achieve promising results. Different from above studies, we focus on exploiting the inter-task correlations among the three aspect-level subtasks and thus incrementally boost one another. Besides, we observe the task characteristics and then use the document-level corpora to pertinently help the corresponding aspect-level subtasks.

Capsule Network. Capsule network (Sabour et al., 2017) has been widely applied in many natural language processing tasks. In ABSA, Wang et al. (2019) focus on building multiple capsules for aspect category sentiment analysis, which do not employ the routing procedure. Chen and Qian (2019) construct a transfer capsule network for transferring semantic knowledge from DSC to SC via sharing the encoder, which utilizes the vanilla capsule network only for the SC subtask. Du et al. (2019) combine capsule network with interactive attention to model the interactive relationship between the given aspect term and context for the SC subtask. Jiang et al. (2019) release a new large-scale multi-aspect multi-sentiment dataset and use capsule network building a strong baseline. Unlike these methods, we pay attention to the end-to-end ABSA task rather than the individual subtask, and propose a dynamic-length to dynamic-length routing algorithm, which can efficiently perform the multi-knowledge transfer.

7 Conclusion

In this paper, we propose an iterative multi-knowledge transfer network for the ABSA task, which can fully exploit the inter-task correlations among the three aspect-level subtasks with the proposed routing algorithm. Moreover, we design a more fine-grained method enabling our model to incorporate the document-level knowledge for pertinently enhancing the corresponding aspect-level tasks. Experimental results on three benchmark datasets demonstrate the effectiveness of our proposed approach, which yields state-of-the-art performance on most metrics.

Acknowledgements

The research work described in this paper has been supported by the National Key R&D Program of China (2019YFB1405200) and the National Nature Science Foundation of China (No. 61976015, 61976016, 61876198 and 61370130). The authors would like to thank the anonymous reviewers for their valuable comments and suggestions to improve this paper.

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. *Tensorflow: A system for large-scale machine learning*. In *USENIX*, pages 265–283.
- Stefanos Angelidis and Mirella Lapata. 2018. *Summarizing opinions: Aspect extraction meets sentiment prediction and they are both weakly supervised*. In *EMNLP*, pages 3675–3686.
- Lingxian Bao, Patrik Lambert, and Toni Badia. 2019. *Attention and lexicon regularized LSTM for aspect-based sentiment analysis*. In *ACL*, pages 253–259.
- Peng Chen, Zhongqian Sun, Lidong Bing, and Wei Yang. 2017. *Recurrent attention network on memory for aspect sentiment analysis*. In *EMNLP*, pages 452–461.
- Zhuang Chen and Tiejun Qian. 2019. *Transfer capsule network for aspect level sentiment classification*. In *ACL*, pages 547–556.
- Zhuang Chen and Tiejun Qian. 2020. *Relation-aware collaborative learning for unified aspect-based sentiment analysis*. In *ACL*, pages 3685–3694.
- Hongliang Dai and Yangqiu Song. 2019. *Neural aspect and opinion term extraction with mined rules as weak supervision*. In *ACL*, pages 5268–5277.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. *BERT: pre-training of deep bidirectional transformers for language understanding*. *CoRR*, abs/1810.04805.
- Li Dong, Furu Wei, Chuanqi Tan, Duyu Tang, Ming Zhou, and Ke Xu. 2014. *Adaptive recursive neural network for target-dependent twitter sentiment classification*. In *ACL*, pages 49–54.
- Chunning Du, Haifeng Sun, Jingyu Wang, Qi Qi, Jianxin Liao, Tong Xu, and Ming Liu. 2019. *Capsule network with interactive attention for aspect-level sentiment classification*. In *EMNLP*, pages 5492–5501.

- Zhifang Fan, Zhen Wu, Xin-Yu Dai, Shujian Huang, and Jiajun Chen. 2019. [Target-oriented opinion words extraction with target-fused neural sequence labeling](#). In *NAACL*, pages 2509–2518.
- Xavier Glorot and Yoshua Bengio. 2010. [Understanding the difficulty of training deep feedforward neural networks](#). In *AISTATS*, pages 249–256.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2018. [Exploiting document knowledge for aspect-level sentiment classification](#). In *ACL*, pages 579–585.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2019. [An interactive multi-task learning network for end-to-end aspect-based sentiment analysis](#). In *ACL*, pages 504–515.
- Mengting Hu, Shiwan Zhao, Li Zhang, Keke Cai, Zhong Su, Renhong Cheng, and Xiaowei Shen. 2019a. [CAN: Constrained attention networks for multi-aspect sentiment analysis](#). In *EMNLP*, pages 4593–4602.
- Minghao Hu, Yuxing Peng, Zhen Huang, Dongsheng Li, and Yiwei Lv. 2019b. [Open-domain targeted sentiment analysis via span-based extraction and classification](#). In *ACL*, pages 537–546.
- Qingnan Jiang, Lei Chen, Ruifeng Xu, Xiang Ao, and Min Yang. 2019. [A challenge dataset and effective models for aspect-based sentiment analysis](#). In *EMNLP*, pages 6280–6285.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. [A convolutional neural network for modelling sentences](#). In *ACL*, pages 655–665.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *EMNLP*, pages 1746–1751.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *EMNLP*, pages 388–395.
- Hao Li and Wei Lu. 2017. [Learning latent sentiment scopes for entity-level sentiment analysis](#). In *AAAI*.
- Xin Li, Lidong Bing, Wai Lam, and Bei Shi. 2018a. [Transformation networks for target-oriented sentiment classification](#). In *ACL*, pages 946–956.
- Xin Li, Lidong Bing, Piji Li, and Wai Lam. 2019a. [A unified model for opinion target extraction and target sentiment prediction](#). In *AAAI*, pages 6714–6721.
- Xin Li, Lidong Bing, Piji Li, Wai Lam, and Zhimou Yang. 2018b. [Aspect term extraction with history attention and selective transformation](#). In *IJCAI*, pages 4194–4200.
- Xin Li, Lidong Bing, Wenxuan Zhang, and Wai Lam. 2019b. [Exploiting BERT for end-to-end aspect-based sentiment analysis](#). In *EMNLP*, pages 34–41.
- Xin Li and Wai Lam. 2017. [Deep multi-task learning for aspect term extraction with memory interaction](#). In *EMNLP*, pages 2886–2892.
- Yunlong Liang, Fandong Meng, Jinchao Zhang, Yufeng Chen, Jinan Xu, and Jie Zhou. 2021. [A dependency syntactic knowledge augmented interactive architecture for end-to-end aspect-based sentiment analysis](#). *Neurocomputing*, 454:291–302.
- Yunlong Liang, Fandong Meng, Jinchao Zhang, Jinan Xu, Yufeng Chen, and Jie Zhou. 2019. [A novel aspect-guided deep transition model for aspect based sentiment analysis](#). In *EMNLP*, pages 5572–5584.
- Peiqin Lin and Meng Yang. 2020. [A shared-private representation model with coarse-to-fine extraction for target sentiment analysis](#). In *Findings of the EMNLP*, pages 4280–4289, Online.
- Kang Liu, Liheng Xu, and Jun Zhao. 2013. [Syntactic patterns versus word alignment: Extracting opinion targets from online reviews](#). In *ACL*, pages 1754–1763.
- Kang Liu, Liheng Xu, and Jun Zhao. 2014. [Extracting opinion targets and opinion words from online reviews with graph co-ranking](#). In *ACL*, pages 314–324.
- Pengfei Liu, Shafiq Joty, and Helen Meng. 2015. [Fine-grained opinion mining with recurrent neural networks and word embeddings](#). In *EMNLP*, pages 1433–1443.
- Huaishao Luo, Tianrui Li, Bing Liu, and Junbo Zhang. 2019. [DOER: Dual cross-shared RNN for aspect term-polarity co-extraction](#). In *ACL*, pages 591–601.
- Dehong Ma, Sujian Li, Fangzhao Wu, Xing Xie, and Houfeng Wang. 2019. [Exploring sequence-to-sequence learning in aspect term extraction](#). In *ACL*, pages 3538–3547.
- Yukun Ma, Haiyun Peng, and Erik Cambria. 2018. [Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive lstm](#). In *AAAI*.
- Margaret Mitchell, Jacqui Aguilar, Theresa Wilson, and Benjamin Van Durme. 2013. [Open domain targeted sentiment](#). In *EMNLP*, pages 1643–1654.
- Thien Hai Nguyen and Kiyooki Shirai. 2015. [PhraseRNN: Phrase recursive neural network for aspect-based sentiment analysis](#). In *EMNLP*, pages 2509–2514.
- Haiyun Peng, Lu Xu, Lidong Bing, Fei Huang, Wei Lu, and Luo Si. 2020. [Knowing what, how and why: A near complete solution for aspect-based sentiment analysis](#). *the AAAI*, 34:8600–8607.

- Minlong Peng, Qi Zhang, Yu-gang Jiang, and Xuan-jing Huang. 2018. [Cross-domain sentiment classification with target domain specific information](#). In *ACL*, pages 2505–2513.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *EMNLP*, pages 1532–1543.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. [SemEval-2015 task 12: Aspect based sentiment analysis](#). In *SemEval 2015*, pages 486–495.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [Semeval-2014 task 4: Aspect based sentiment analysis](#). In *SemEval 2014*, pages 27–35.
- Guang Qiu, Bing Liu, Jiajun Bu, and Chun Chen. 2011. Opinion word expansion and target extraction through double propagation. *Computational linguistics*, 37(1):9–27.
- Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. 2017. Dynamic routing between capsules. In *NIPS*, pages 3856–3866.
- Martin Schmitt, Simon Steinheber, Konrad Schreiber, and Benjamin Roth. 2018. [Joint aspect and polarity classification for aspect-based sentiment analysis with end-to-end neural networks](#). In *EMNLP*, pages 1109–1114.
- Kai Sun, Richong Zhang, Samuel Mensah, Yongyi Mao, and Xudong Liu. 2019. [Aspect-level sentiment analysis via convolution over dependency tree](#). In *EMNLP*, pages 5683–5692.
- Duyu Tang, Bing Qin, Xiaocheng Feng, and Ting Liu. 2016. [Effective lstms for target-dependent sentiment classification](#). In *COLING*, pages 3298–3307.
- Jialong Tang, Ziyao Lu, Jinsong Su, Yubin Ge, Linfeng Song, Le Sun, and Jiebo Luo. 2019. [Progressive self-supervised attention learning for aspect-level sentiment analysis](#). In *ACL*, pages 557–566.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *NIPS*, pages 5998–6008.
- Duy-Tin Vo and Yue Zhang. 2015. Target-dependent twitter sentiment classification with rich automatic features. In *IJCAI*.
- F. Wang, M. Lan, and W. Wang. 2018. [Towards a one-stop solution to both aspect extraction and sentiment analysis tasks with neural multi-task learning](#). In *IJCNN*, pages 1–8.
- Shuai Wang, Sahisnu Mazumder, Bing Liu, Mianwei Zhou, and Yi Chang. 2018. [Target-sensitive memory networks for aspect sentiment classification](#). In *ACL*, pages 957–967.
- Wenya Wang, Sinno Jialin Pan, Daniel Dahlmeier, and Xiaokui Xiao. 2016a. [Recursive neural conditional random fields for aspect-based sentiment analysis](#). In *EMNLP*, pages 616–626.
- Wenya Wang, Sinno Jialin Pan, Daniel Dahlmeier, and Xiaokui Xiao. 2017. Coupled multi-layer attentions for co-extraction of aspect and opinion terms. In *AAAI*.
- Yequan Wang, Minlie Huang, xiaoyan zhu, and Li Zhao. 2016b. [Attention-based lstm for aspect-level sentiment classification](#). In *EMNLP*, pages 606–615.
- Yequan Wang, Aixin Sun, Minlie Huang, and Xiaoyan Zhu. 2019. [Aspect-level sentiment analysis using capsules](#). In *WWW*, pages 2033–2044.
- Hu Xu, Bing Liu, Lei Shu, and Philip Yu. 2020a. [DomBERT: Domain-oriented language model for aspect-based sentiment analysis](#). In *Findings of the EMNLP*, pages 1725–1731.
- Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2018. [Double embeddings and CNN-based sequence labeling for aspect extraction](#). In *ACL*, pages 592–598.
- Lu Xu, Hao Li, Wei Lu, and Lidong Bing. 2020b. [Position-aware tagging for aspect sentiment triplet extraction](#). In *the EMNLP*, pages 2339–2349.
- Yichun Yin, Furu Wei, Li Dong, Kaimeng Xu, Ming Zhang, and Ming Zhou. 2016. [Unsupervised word and dependency path embeddings for aspect term extraction](#). *CoRR*, abs/1605.07843.
- Meishan Zhang, Yue Zhang, and Duy-Tin Vo. 2015. [Neural networks for open domain targeted sentiment](#). In *EMNLP*, pages 612–621.
- Yan Zhou, Longtao Huang, Tao Guo, Jizhong Han, and Songlin Hu. 2019. [A span-based joint model for opinion target extraction and target sentiment classification](#). In *IJCAI*, pages 5485–5491.

Appendix

A Implementation Details

Following (Chen and Qian, 2020), we use 300d GloVe released by Pennington et al. (2014) as general-specific embeddings and the embeddings released by Xu et al. (2018) as domain-specific embeddings. Our models are trained by Adam optimizer (Kingma and Ba, 2014), with learning rate $\eta_0 = 10^{-4}$, and batch size is set to 32. When training, we randomly sample 20% of each training data as the validation set and the remaining 80% as training set.

We following (Chen and Qian, 2020) fix the domain-specific and general-specific word embeddings in all models, where the domain-specific embedding vectors are 100 dimensions. The trainable weight matrices in the *CNN* are initialized by following the Glorot Uniform strategy (Glorot and Bengio, 2010). Besides, all biases are initialized as zero. We tune the number of *CNN* layer on the validation set of each dataset. Finally, The *CNN* layer number in the shared encoder is set to 2, and is fixed as 2, 2, 1 for the ATE subtask, the OTE subtask, and the ASC subtask in task-specific layers, respectively. The *CNN* layer in the shared encoder has 150 filters with kernel size $k = 3$ and 150 filters with kernel size $k = 5$. The *CNN* layers in each task-specific encoder have 300 filters with kernel size $k = 5$ per layer. The activation function is ReLU for each *CNN* layer. Dropout is employed after the embedding layer and each *CNN* layer, which is empirically set to 0.5. The discount coefficients $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ and λ_5 in loss functions are not fine-tuned and empirically set to 1.0.

Since the extracted aspect term may consist of several tokens and the predicted polarity of each token may be inconsistent, we thus following (Chen and Qian, 2020) only take the sentiment polarity of the first token of the current aspect term as the sentiment label for measuring the performance. We also note that only aspect terms have sentiment annotations and thus following (Chen and Qian, 2020) only consider ASC predictions on these aspect term-related tokens for computing the ASC loss and ignore the sentiments predicted on other tokens.

For training, we first train the model with document-level tasks for five epochs, and then alternately train our model on aspect-level tasks with 2 epochs and document-level tasks with 1 epoch. Finally, we train the model for a fixed number of epochs, and obtain the best results at the epoch with the best F1-absa score on the validation set for producing the testing results, as did in (He et al., 2019).

In our experiments, following (Chen and Qian, 2020), we also use *BERT_{LARGE}* (Devlin et al., 2018) as the backbone to further investigate our model performance.

The neural model is implemented in Keras and all computations are done on an NVIDIA Tesla V100 GPU, where each experiment runs about 1~3 hours. Hyperparameter configurations for best-

performing models have explained above. The method of choosing hyperparameter values is manual tuning on the validation and the criterion used to select is F1-absa. The downloadable version of used data can be found in: <https://github.com/ruidan/IMN-E2E-ABSA>, provided by IMN (He et al., 2019), where we use this data without any pre-processing.

B Experiments of Hyperparameters

Impact of Iteration Number: T .

As an important hyperparameter, we investigate the impact of iterations T . Tab. 7 shows the change of F1-absa on the validation set of each dataset. We find that the best results can be obtained when T equals 1, 2, and 4, respectively. There is no consistent conclusion about how to set this parameter. In general, T is set to 1, 2, and 4 on D1, D2, and D3 in our experiments, respectively.

T	0	1	2	3	4	5
D1	62.78	63.56	63.14	63.44	63.00	62.34
D2	53.34	55.25	56.22	56.07	55.47	54.88
D3	65.04	65.72	65.88	65.72	66.35	65.78

Table 7: F1-absa (%) scores with different T values. Average results over 5 runs on the validation set are reported.

Impact of Routing Number: $iter$.

Tab. 8 (in the next page) shows the impact of the maximum number of the routing number $iter$ of the routing algorithm on the validation set of each dataset. The results demonstrate that the model achieves the best results when routing number equals 3 and further iterations do not further improve the performance. In general, the routing number is fixed to 3 in our experiments.

$iter$	1	2	3	4	5
D1	63.06	63.80	64.52	64.02	64.25
D2	56.28	56.47	57.14	56.70	56.47
D3	65.71	66.32	66.75	66.03	66.00

Table 8: F1-absa (%) scores with different routing number in Routing Block. Average results over 5 runs on the validation set are reported.

C Error Analysis

We have checked some error examples and made a thorough error analysis, which can be roughly divided into 3 types. 1) Due to aspect extraction and

Sentence	The service is slow.
Aspect	service
Opinion	slow
Sentiment Polarity	negative
Aspect-Sentiment Pair	service-negative
Aspect-Opinion Pair	service-slow
Aspect-Opinion-Sentiment Triplet	service-slow-negative

Table 9: The example of Aspect-Sentiment Pair, Aspect-Opinion Pair, and Aspect-Opinion-Sentiment Triplet.

opinion extraction are not always correctly identified, the Aspect-Opinion-Sentiment triplet is hard to handle. 2) The imbalanced label distribution in the training corpus. 3) The complex instances are hard to correctly deal with, such as the sentence that has multiple aspects and opinions, which are hardly effectively learned. For instance, in the sentence “coffee is a better deal than overpriced cosi sandwiches”, where two opinion terms “better” and “overpriced”, and two aspect terms “coffee” and “cosi sandwiches” are mentioned, where the sentiment polarities of them are “positive” and “negative”, respectively. In this case, our IMKTN correctly extracted all aspect terms, and the IMKTN successfully detected the opinion term “better” but failed to identify the opinion term “overpriced”, *i.e.*, the OTE subtask failed partly, where the IMKTN made right sentiment classification for the aspect term “coffee” but assigned wrong sentiment polarity (“positive”) to the aspect term “cosi sandwiches”. The reason may be that the knowledge from the opinion term “better” contributed to the right sentiment classification for “coffee” but led to the wrong sentiment classification for “cosi sandwiches”. If the opinion term “overpriced” can be successfully identified, it may contribute to the right classification for “cosi sandwiches” with our routing algorithm.