

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/139527>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Aspect-Invariant Sentiment Feature Learning: Adversarial Multi-task Learning for Aspect-Based Sentiment Analysis

ABSTRACT

Neural models with attention mechanisms have achieved remarkable performance in Aspect-based Sentiment Analysis (ABSA). In most previous studies, the information about aspects in sentences is considered important for the ABSA task and therefore various attention mechanisms have been explored to leverage interactions between aspects and context. However, some sentiment expressions carry the same polarity regardless of the aspects they are associated with. In such cases, it is not necessary to incorporate aspect information for ABSA. In fact, in our experiments, we find that blindly leveraging interactions between aspects and context as features may introduce noise when analyzing those aspect-invariant sentiment expressions, especially when facing with limited aspect-related annotated data. Hence, in this paper, we propose an Adversarial Multi-task Learning Framework to identify the aspect-invariant/dependent sentiment expressions automatically without requiring extra annotations. In addition, we use a gating mechanism to control the contribution of representations derived from aspect-invariant and aspect-dependent hidden states when generating the final contextual sentiment representations for the given aspect. This essentially allows the exploitation of aspect-invariant sentiment features for better ABSA results. Experimental results on two benchmark datasets show that extending existing neural models using our proposed framework achieves superior performance. In addition, the aspect-invariant data extracted by our framework can be considered as pivot features for better transfer learning of the ABSA models on unseen aspects.

KEYWORDS

Aspect sentiment analysis, Aspect invariant sentiment, Adversarial training, Multi-task learning

ACM Reference Format:

. 2020. Aspect-Invariant Sentiment Feature Learning: Adversarial Multi-task Learning for Aspect-Based Sentiment Analysis. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Aspect-based sentiment analysis (ABSA) aims at inferring the sentiment polarity of a specific aspect expressed in a sentence [20, 21]¹. For example, in the sentence *'The food is good, but the service is*

terrible', there are two aspects mentioned: *'FOOD#QUALITY'* (corresponding to the aspect entity *'food'*) and *'SERVICE#GENERAL'* (corresponding to the aspect entity *'service'*). Here, each aspect consists of an entity and an attribute. The sentiment polarity of the aspect *'FOOD#QUALITY'* is positive while the sentiment polarity of the aspect *'SERVICE#GENERAL'* is negative.

Thanks in part to the availability of copious annotated resources for some aspects, recent attention-based models can effectively distinguish the sentiment polarities of different aspects in the same sentence [9, 18, 41]. Despite remarkable progress made in ABSA, most existing work only focused on how to extract contextual sentiment information for specific aspects [8, 23, 27, 32, 45], or learn aspect-dependent features for sentiment classification [14, 16, 28, 40]. However, in some cases, polarities carried by sentiment expressions are aspect-independent and blindly incorporating the associated aspect information may confuse the sentiment classifier, especially when facing with limited aspect-related annotated data. Hence, we argue that separating aspect-invariant sentiment expressions from aspect-dependent ones could potentially lead to improved ABSA results. To be best of our knowledge, there is no prior work focusing on identifying the aspect-invariant and aspect-dependent sentiment expressions with no supervised information.

To separate aspect-invariant sentiment expressions from aspect-dependent ones, we have the following observations:

- If a sentiment expression is aspect-invariant, we can simply replace its associated aspect entity with another one without inverting the polarity. In this way, it is possible to automatically generate synthetic training instances to augment the training data.
- If a sentiment expression is aspect-dependent, then simply replacing its aspect may change its polarity or derive a noisy sample. We could train a discriminator by adversarial training to identify those synthetic training instances and the original training examples.

To illustrate our idea, we give examples shown in Figure 1 where some sentence examples are paired with their corresponding aspects and polarity labels. In Figure 1(a), suppose there are no training examples for the aspect *'LOCATION#GENERAL'* in the training set, in such cases, many existing sentiment analysis systems may fail to detect the aspect-specific sentiment polarity of *'Oh yeah the view was good too.'* enclosed in the green box. However, there are training examples for other aspects such as *'SERVICE#GENERAL'* which share a similar sentiment expression as shown in the sentence enclosed by the red box. We could generate a synthetic training example enclosed in the blue box by replacing the word *'service'* with the word *'location'*. With the augmented training data, models could better capture such aspect-invariant sentiment features and thus generate better ABSA results. However, for sentence containing aspect-dependent sentiment expressions, we cannot generate synthetic training examples by simply replacing the aspect entity

¹This work focuses on sentiment analysis of aspect categories.

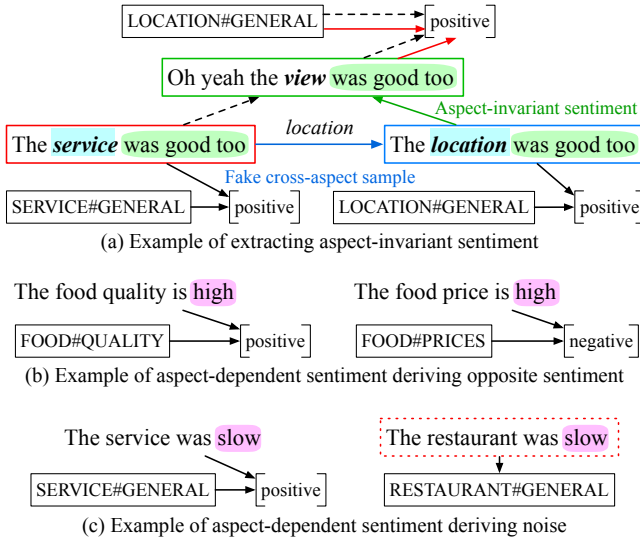


Figure 1: Sentence examples paired with their aspects and polarity labels. (a) All three sentences share the same sentiment expression. The sentence enclosed in the red box is seen in the training set, while the sentence enclosed in the blue box is the synthetic training example generated by replacing the aspect entities from ‘service’ to ‘location’. The created synthetic example would allow the detection of the polarity of the unseen test example enclosed in the green box; (b) Examples of aspect-dependent sentiment expressions. Although these two sentences share the same sentiment expression, they express opposing polarities. (c) Examples of another type of aspect-dependent sentiment expressions, in which, the word ‘slow’ can be used to modify ‘service’ but not ‘restaurant’.

or attribute with other aspect entities or attributes as the resulting polarity may be inverted. In Figure 1(b), the sentiment word ‘high’ is positive for the aspect ‘FOOD#QUALITY’ but negative for ‘FOOD#PRICES’. In some other instances, simply adopting aspect-dependent sentiment expressions to generate synthetic training examples may derive noise, such as the sentence enclosed by the dotted red box in Figure 1(c). Hence, separating the aspect-invariant sentiment features from aspect-dependent ones should be considered for improving the performance of ABSA.

In this paper, we propose an adversarial multi-task learning framework to extract aspect-invariant sentiment features, and distinguish aspect-invariant sentiment expressions from aspect-dependent ones by adversarial training. We first generate fake cross-aspect samples by replacing the aspect entities in the original training samples with other aspect entities but keeping the original polarity labels unchanged. If the original training sample contains aspect-invariant sentiment expressions, then the polarity label of a synthetic example should be correct and the discriminator would not be able to distinguish between the real training example and the fake example. If the original sample contains aspect-dependent sentiment expressions, then it is likely that the polarity label of the fake example would be wrong and the discriminator can easily identify the fake example. In this way, we can use adversarial training to obtain a large number of

aspect-invariant and aspect-dependent training instances. For different types of sentiment expressions, the representation composition is controlled by a discriminator-based gate to capture sentiment features better by judging if the aspect-based representation is needed or not. Hence, the ABSA performance can be improved by augmenting training examples containing aspect-invariant sentiment features, especially for those aspects with limited training examples in the training set. In addition, aspect-dependent features are not blindly incorporated during representation learning and will only be added if necessary. The main contributions of our work can be summarized as follows:

- The ABSA task is approached from a new perspective that aspect-invariant sentiment features are leveraged for sentiment analysis, especially for those aspects with limited training examples.
- A novel multi-task discriminator is proposed for learning aspect-invariant/dependent sentiment features in an adversarial way.
- As a general adversarial multi-task learning framework, our proposed method can be easily combined with any ABSA models or other neural networks to generate improved results.

To the best of our knowledge, our work represents the first study of extracting aspect-invariant and aspect-dependent sentiment features by adversarial learning for ABSA.

2 RELATED WORK

2.1 Aspect-based Sentiment Analysis

The task of aspect-based sentiment analysis (ABSA) can be regarded as a fine-grained sentiment analysis task, which needs to leverage information from both context and aspects [3, 5–7, 11, 17, 35–37, 44]. Xue and Li [44] proposed a gated CNN model to selectively output the sentiment features according to the given aspect based on gating units. Chen et al. [1] proposed a memory network-based framework with multiple-attention mechanism to capture sentiment features separated by a long distance, so that it is more robust against irrelevant information. Tang et al. [37] utilized two LSTMs to extract the contextual sentiment dependencies for a given aspect. Resoundingly, thanks to the ability of attention mechanisms, the majority of current approaches attempt to enforce models to pay more attentions to the given aspect in the training process [1, 9, 10, 12, 18, 25, 38, 41]. For example, an attention-based LSTM by means of using an individual aspect embedding is able to focus on the aspect and capture the intra-sentence relations [41]. There is also an interactive attention network with context representation and aspect representation learned separately and interactively to well represent the context and its corresponding target [25]. In addition, with the development of graphical neural networks, an aspect-specific graph convolutional network (GCN) based on dependency trees was proposed to learn contextual dependencies for a given aspect [46]. Recently, based on the remarkable success of BERT [4], Song et al. [34] proposed an attentional encoder network to draw hidden states and semantic interactions between target and context words. Sun et al. [35] constructed an auxiliary sentence from an aspect and convert ABSA to a sentence-pair classification task based on fine-tuning the pre-trained model from BERT. In addition, Li et al. [19] exploited coarse-to-fine task transfer and proposed a multi-granularity alignment network

for simultaneously aligning aspect granularity and aspect-specific feature representations across domains. This work authentically borrowed knowledge from an abundant source domain of the coarse-grained aspect category task to a small-scale target domain of the fine-grained aspect term task. However, it did not address the problem of limited annotated data for some aspects in the training data.

Almost all the aforementioned models assume contextual sentiment information towards a specific aspect is essential for ABSA, which is however not always the case. Blindly incorporating aspect-related sentiment features may lead to degraded sentiment classification performance, as will shown in our experiments. To address this issue, a better network for leveraging aspect-invariant sentiment information should be considered.

2.2 Adversarial Multi-task Learning

By extracting shared and transferable information from related tasks, multi-task learning can leverage latent correlated features from data [24, 33, 42]. In the task of ABSA, He et al. [9] considered ABSA as two subtasks: document-level and aspect-level classification, and transfer knowledge from document-level data to improve the performance of aspect-level sentiment classification. Further, He et al. [10] utilized an iterative message passing scheme to explicitly model the interactions between tasks. Multi-task learning could be combined with adversarial training, which has achieved promising performance in various natural language processing (NLP) tasks [2, 15, 22, 39, 43]. Among them, Wu et al. [43] applied adversarial training in relation extraction within the multi-instance multi-label learning framework, which revealed the effectiveness of adversarial training for relation extraction. Liu et al. [22] proposed an adversarial multi-task learning framework for text classification, in which the shared and private feature spaces are inherently disjoint by introducing orthogonality constraints.

In sentiment analysis, Wang et al. [39] proposed a user-attention-based CNN model with adversarial cross-lingual learning framework to enrich the user post representation in personalized Microblog sentiment classification, and Chen et al. [2] proposed an adversarial deep averaging network to transfer the knowledge learned from labeled data on a resource-rich source language to low-resource languages in the task of cross-lingual sentiment classification. Li et al. [18] considered ABSA as an end-to-end task, and proposed a selective adversarial learning method to learn an alignment weight for each word, where more important words can possess higher alignment weights to achieve a local semantic alignment and capture domain-invariant word representations by employing adversarial training. Inspired by the recent successful work [2, 22], we propose a novel framework to extract aspect-invariant sentiment features for improving aspect-based sentiment classification.

3 METHODOLOGY

In this section, we describe our adversarial multi-task learning framework, which is demonstrated in Figure 2, in details.

3.1 Fake Cross-Aspect Sample Generation

Ideally, we would expect a sufficient number of training examples annotated for each aspect in the task of ABSA. However, this is rarely the case in reality. In existing ABSA datasets, we often observe

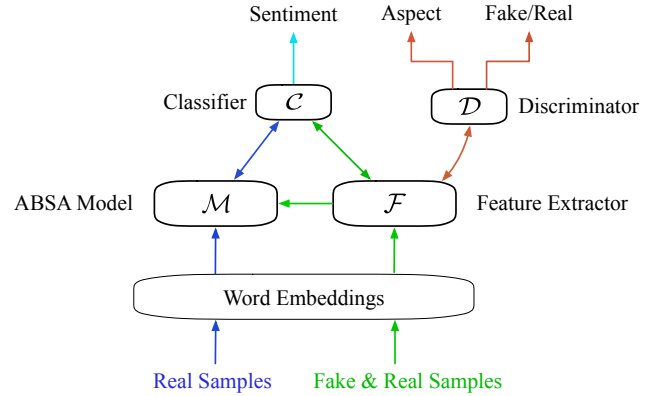


Figure 2: The architecture of the proposed adversarial multi-task learning framework. There are five main components: the embedding layer, the feature extractor \mathcal{F} , the discriminator \mathcal{D} , the ABSA model \mathcal{M} and the sentiment classifier \mathcal{C} . Lines in different colors show the propagation of information from different components.

imbalanced training examples with some aspects associated with abundant annotated data while others having only limited annotated instances. The goal of fake cross-aspect sample generation is to automatically generate high-quality training samples so as to enrich the training dataset.

There are many possible ways to generate synthetic fake training examples. One simple approach we consider is to replace an aspect mention or attribute in a sentence with another aspect term which shares some similarity with the original aspect mention/attribute². The rationale behind this is that we expect sentiment expressions of these related aspects are also similar. For example, as shown in Figure 1(a), to generate the fake training example for the aspect ‘LOCATION#GENERAL’, we select the training instances which are annotated with the attribute ‘#GENERAL’ and replace its corresponding entity mention with the aspect entity ‘location’.

3.2 Adversarial Multi-Task Learning Framework

As demonstrated in Figure 2, our adversarial multi-task learning framework consists of two branches: aspect-based sentiment prediction and aspect-invariant features extraction. As a universally applicable framework, the proposed framework can be easily combined with existing ABSA models³ directly.

We first assume a sentence is represented by a sequence of words $\{w_1, w_2, \dots, w_n\}$, and it may contain one or more aspects associated with different sentiment polarities, i.e. *positive*, *negative*, and *neutral*. Each word in the sentence can be represented by an m -dimensional embedding $\mathbf{x} \in \mathbb{R}^k$, where k is the dimension of word embeddings. For each input sentence, the embedding layer maps the corresponding embedding vector for each word from the full word embedding matrix $\mathbf{V} \in \mathbb{R}^{k \times |L|}$ to get the input embedding matrix

²In our experiments, we ignore the training examples whose entities were annotated with ‘NULL’, since they are unable to locate a specific entity and only accounted for a small part of the training set.

³Such as ATAE-LSTM, RAM, GCAE etc., which will be discussed in the Experiments section.

$\mathbf{X} \in \mathbb{R}^{k \times n}$, where n is the length of sentence and $|L|$ is the vocabulary size of the dataset. The embedding layer is usually initialized with pre-trained embeddings such as GloVe [29] and BERT [4], then they are fine-tuned during the training process. Each aspect can also be represented as an m -dimensional embedding $\mathbf{a} \in \mathbb{R}^k$, which is the average of its entity and attribute word embeddings.

After that, a feature extractor \mathcal{F} is utilized to extract the aspect-invariant sentiment features, and meanwhile separate them from aspect-dependent sentiment features and noisy features. Here, the feature extractor can be any neural model, such as bi-directional LSTM (Bi-LSTM):

$$\mathbf{h}_t^f = \mathcal{F}(\mathbf{X}_t) = \text{Bi-LSTM}(\mathbf{X}_t, \mathbf{h}_{t-1}^f, \theta^f), \quad (1)$$

where $\text{Bi-LSTM}(\cdot)$ is the shorthand computation of Bi-LSTM, $\mathbf{X}_t \in \mathbb{R}^k$ is the input of word embedding at the current time step, \mathbf{h}_{t-1}^f is the hidden output at the last time step, θ^f denotes all the parameters of the feature extractor \mathcal{F} . Then the hidden output of feature extractor $\mathcal{F}(\mathbf{X}) = \mathbf{h}^f$ can be fed as into the discriminator \mathcal{D} to identify the aspect label and the sample label (*fake* or *real*). For sentiment prediction of fake cross-aspect samples, we feed $\mathcal{F}(\mathbf{X})$ into the sentiment classifier \mathcal{C} to get the output distribution. Here, we utilize the *softmax* function to obtain the output distribution of the discriminator and the classifier.

For the branch of sentiment prediction for a specific aspect, the input embedding matrix $\mathbf{X} \in \mathbb{R}^{k \times n}$ and the aspect embedding $\mathbf{a} \in \mathbb{R}^k$ are fed into the ABSA model \mathcal{M} and output the hidden feature representation \mathbf{h}^m :

$$\mathbf{h}^m = \mathcal{M}(\mathbf{X}) = \text{ABSAModel}(\mathbf{X}, \mathbf{a}, \theta^m), \quad (2)$$

where θ^m represents the parameters of the ABSA model. To capture the aspect-invariant sentiment features, we feed the hidden outputs from the ABSA model \mathcal{M} and the feature extractor \mathcal{F} into the sentiment classifier \mathcal{C} to obtain the sentiment polarity of a specific aspect:

$$\mathbf{y}^c = \mathcal{C}(\mathbf{X}) = \text{softmax}(\mathbf{W} \cdot [\mathbf{h}^f, \mathbf{h}^m] + \mathbf{b}), \quad (3)$$

where \mathbf{y}^c is the predicted sentiment distribution. Here, we also utilize the *softmax* function for sentiment prediction. $[\cdot, \cdot]$ represents a concatenation of two feature representations.

3.3 Multi-Task Learning for ABSA

The key purpose of multi-task learning is to share latent features extracted from related tasks [2, 10]. Here, the feature extractor \mathcal{F} aims at reconciling the features learned from different tasks with the help of the multi-task discriminator \mathcal{D} .

Aspect Discrimination. The sentiment expressed in a sentence might depend on a specific aspect. Here, a discriminator \mathcal{D} is adopted to predict the aspect category given a sentence input by minimizing the cross-entropy loss of predicted and true aspect distributions for all N training samples:

$$\mathcal{J}^a(\theta_d) = - \sum_{i=1}^N \sum_{j=1}^A y_i^j \log \mathcal{D}(\mathcal{F}(\mathbf{X})), \quad (4)$$

where A is the number of aspect classes, y_i^j is the ground-truth distribution of aspect, θ_d represents all the parameters of \mathcal{D} .

Algorithm 1 The training procedure of our adversarial framework

Require: embedding matrix sets of samples \mathbb{X} ; aspect embedding sets \mathbb{A} ; the size of dataset N , N_r ; hyperparameters α , β and λ .

```

1: for  $i = 1$  to  $N$  do
2:    $\mathbf{X} = \mathbb{X}_i$ 
3:    $\mathbf{a} = \mathbb{A}_i$ 
4:    $\triangleright$  Aspect and fake sample discrimination
5:    $\mathbf{f}_{\text{cross}} = \mathcal{F}(\mathbf{X})$ 
6:    $\mathcal{J}_d = \text{CrossEntropy}(\mathcal{D}(\mathbf{f}_{\text{cross}}))$ 
7:    $\triangleright$  Sentiment discrimination
8:    $\mathcal{J}_{sc} = \text{CrossEntropy}(\mathcal{C}(\mathbf{f}_{\text{cross}})) - \alpha \mathcal{J}_d$ 
9:    $\triangleright$  Aspect-based sentiment prediction
10:   $\mathbf{f}_m = \mathcal{M}(\mathbf{X}, \mathbf{a})$ 
11:   $\mathcal{J}_{ac} = \text{CrossEntropy}(\mathcal{C}(\mathbf{f}_m)) + \lambda \mathcal{J}_{sc}$ 
12: end for
13: Update all the parameters to minimize the loss

```

Fake Sample Discrimination. To capture aspect-invariant sentiment features from the generated cross-aspect samples, the discriminator \mathcal{D} is also adopted to classify whether a given input sample is real or fake by minimizing the cross-entropy loss of predicted and true distributions:

$$\begin{aligned} \mathcal{J}^s(\theta_d) = & - \sum_{i=1}^N z_i^j \log \mathcal{D}(\mathcal{F}(\mathbf{X})) \\ & + (1 - z_i^j) \log(1 - \mathcal{D}(\mathcal{F}(\mathbf{X}))), \end{aligned} \quad (5)$$

where z_i^j is the ground-truth distribution of fake or real sample.

Sentiment Discrimination. For sentiment discrimination of cross-aspect samples, the sentiment features extracted by \mathcal{F} are fed into classifier \mathcal{C} to extract aspect-invariant sentiment information and discriminate noisy and aspect-dependent sentiment features. The objective to train the classifier \mathcal{C} is defined as minimizing the cross-entropy loss of prediction and true distributions:

$$\mathcal{J}(\theta_{sc}) = - \sum_{i=1}^N \sum_{j=1}^S p_i^j \log \mathcal{C}(\mathcal{F}(\mathbf{X})), \quad (6)$$

where S is the number of sentiment classes, p_i^j is the ground-truth distribution of sentence-level sentiment polarity, θ_{sc} represents all the parameters of sentence-level sentiment discrimination.

Aspect-based Sentiment Prediction. For aspect-based sentiment prediction, the aspect-invariant sentiment features extracted by \mathcal{F} are fed into existing ABSA models. With the help of the aspect-invariant features, the ABSA model \mathcal{M} can learn better sentiment features for those aspects with limited training samples. The objective function is defined as:

$$\mathcal{J}(\theta_{ac}) = - \sum_{i=1}^{N_r} \sum_{j=1}^S q_i^j \log \mathcal{C}(\mathcal{M}(\mathbf{X})), \quad (7)$$

where N_r is the number of real samples, q_i^j is the ground-truth distribution of aspect-based sentiment polarity, θ_{ac} represents all the parameters of aspect-based sentiment prediction.

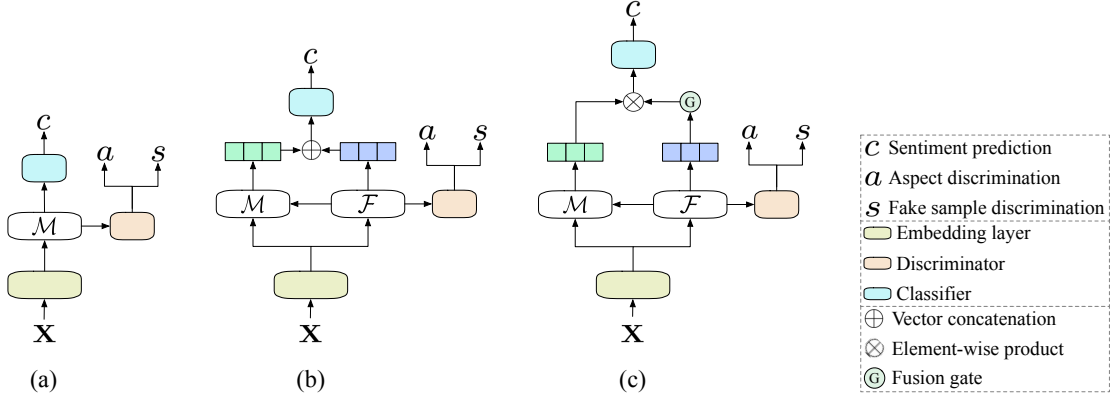


Figure 3: Three schemes of our adversarial multi-task learning framework.

3.4 Adversarial Training

The training procedure of our adversarial framework is depicted in Algorithm 1. The feature extractor \mathcal{F} aims to extract aspect-invariant sentiment features that could help an ABSA model \mathcal{M} to predict sentiment polarity for any aspect and confuse the discriminator \mathcal{D} to detect aspect label and fake sample by adversarial training. In another words, if a well-trained discriminator \mathcal{D} can not detect the aspect label and fake sample based on the features learned by \mathcal{F} , then those features are aspect-invariant and could be exploited to improve aspect-based sentiment prediction.

In pursuit of the adversarial goal, the discriminator \mathcal{D} is designed to impede the feature extractor \mathcal{F} to learn aspect-invariant sentiment features from input sentences. In addition, the sentiment classifier \mathcal{C} needs to predict the sentiment polarity of the input sentence. Here, adversarial training performs min-max optimization that can be divided into two parts: minimizing the cross-entropy loss of sentence-level sentiment prediction and maximizing the cross-entropy loss of the discriminator. Hence, the adversarial loss can be defined as:

$$\mathcal{J}(\Theta_{sc}) = \mathcal{J}(\theta_{sc}) - \alpha \mathcal{J}(\theta_d), \quad (8)$$

where α is the weight that controls the interaction of the loss terms, and $\mathcal{J}(\theta_d)$ denotes the overall loss of the discriminator \mathcal{D} , which is defined as:

$$\mathcal{J}(\theta_d) = \mathcal{J}^s(\theta_d) + \beta \mathcal{J}^a(\theta_d) \quad (9)$$

Finally, we simultaneously minimize the cross-entropy loss of sentence-level sentiment prediction and aspect-based sentiment prediction to exploit the aspect-invariant sentiment features:

$$\mathcal{J}(\Theta) = \mathcal{J}(\theta_{ac}) + \lambda \mathcal{J}(\Theta_{sc}), \quad (10)$$

where Θ represents all the parameters of the adversarial multi-task learning framework, λ is the parameter to control the influence of $\mathcal{J}(\theta_d)$.

3.5 Three Schemes of Our Framework

To demonstrate different scenarios in deploying our proposed framework, we explore three different structures to extract and transfer aspect-invariant sentiment features for improving the performance of aspect-based sentiment prediction, as shown in Figure 3.

Integration structure (ISF). In the first scheme, we integrate the feature extractor \mathcal{F} with an ABSA model \mathcal{M} , i.e. \mathcal{M} is also used as \mathcal{F} . As demonstrated in Figure 3(a), the aspect-invariant sentiment information and aspect-based sentiment information are both learned by the ABSA model \mathcal{M} .

Concatenation structure (CSF). As demonstrated in Figure 3(b), a recurrent model Bi-LSTM is used as the feature extractor \mathcal{F} to extract aspect-invariant sentiment features from data, which are subsequently concatenated with the hidden output of the ABSA model \mathcal{M} :

$$\mathbf{h} = \mathbf{h}^f \oplus \mathbf{h}^m, \quad (11)$$

and fed into the sentiment classifier \mathcal{C} .

Gate fusion structure (GSF). The last scheme is a gated fusion structure, as demonstrated in Figure 3(c). Bi-LSTM is also used to learn aspect-invariant sentiment features. Different from CSF, feature combination is performed by a fusion gate:

$$\mathbf{h} = (\mathbf{h}^m \mathbf{W}^m + \mathbf{b}^m) \otimes \sigma(\mathbf{h}^f \mathbf{W}^f + \mathbf{b}^f), \quad (12)$$

where σ denotes the sigmoid function.

4 EXPERIMENTS

4.1 Datasets and Experimental Setting

We conduct experiments on two benchmark datasets of the restaurants domain from Semeval 2015 Task 12 [31] and Semeval 2016 Task 5 [30] (Sem15 and Sem16). The task of ABSA aims at predicting the sentiment polarity (i.e. *positive*, *negative* or *neutral*) for each aspect mentioned in the sample. Each training and test sample in the datasets consists of the review sentence, aspect entity, aspect, and the sentiment polarity towards the aspect⁴. We use GloVe [29] to initialize word embeddings for all non-BERT models, and use the pre-trained uncased BERT-base [4] model for all BERT-based models. The weight coefficients of α , β and λ are 0.5, 0.1 and 0.5 respectively. We randomly initialize \mathbf{W} and \mathbf{b} for all experiments. To reconcile the training of multi-task, we pre-train the discriminator \mathcal{D} for 20 epochs.

⁴We remove sentences containing no aspects. The remaining sentences contain one or more aspects.

4.2 Comparison Models

We compare our proposed adversarial multi-task learning framework (including three schemes, i.e. **ISF**, **CSF** and **GSF**) with 10 baseline models, including BERT:

- **TD-LSTM** [37]: A target-dependent LSTM model, which incorporates aspect information into LSTM.
- **ATAE-LSTM** [41]: An attention-based LSTM model which better takes advantage of aspect information.
- **MemNet** [38]: An attention-based memory network, which applies attention multiple times on the word embeddings for sentence representation to capture the importance of each context word.
- **RAM** [1]: A recurrent attention-based memory network which captures sentiment features separated by a long distance.
- **IAN** [25]: An interactive attention network with context representation and aspect representation learned separately but interactively.
- **AOA** [13]: An attention-based network models aspects and sentences jointly to capture the relations between aspects and context.
- **MGAN** [7]: A multi-grained attention network which captures the word-level interaction between aspects and context.
- **GCAE** [44]: A gated CNN model which effectively controls the flow of sentiment according to the given aspect information.
- **AEN** [34]: An attention encoder network, which eschews recurrence and employs attention-based encoders for the modeling between context and aspect.
- **BERT** [4]: The vanilla pre-trained uncased BERT-base model, which adopts “[CLS] sentence [SEP] aspect [SEP]” as input⁵.
- **ISF+models**: The models based on our integration structure framework.
- **CSF+models**: The models based on our concatenation structure framework.
- **GSF+models**: The models based on our gated fusion structure framework.

4.3 Main Experiment Results

The main experimental results on two benchmark datasets are reported in Table 1. We can draw a conclusion that the ABSA models based on our adversarial multi-task learning framework (all three schemes, i.e. **ISF+models**, **CSF+models** and **GSF+models**) achieve better performance than competitor models for both accuracy and Macro-F1. Among them, the best improvements of accuracy and Macro-F1 are 5.25% (**CSF+AOA**) and 6.32% (**GSF+MemNet**) respectively on sem15. For the Sem16 dataset, our framework improves accuracy by 4.84% (**GSF+RAM**) and Macro-F1 by 8.75% (**GSF+MGAN**). Among the baselines, BERT gives the best results on both Sem15 and Sem16 datasets. Nevertheless, when integrated with our proposed framework, significant improvements of 2.57-3.13% in accuracy and 2.50-3.63% in F1 are observed compared to the vanilla BERT model. The results verify that our proposed adversarial multi-task learning framework can be easily combined

⁵We have also tried BERT-large in preliminary experiments, and found that BERT-large integrated with our framework also achieves state-of-the-art performance for aspect-based sentiment classification on two benchmark datasets.

Table 1: Main experimental results on Sem15 and Sem16. Acc. represents accuracy, F1 represents Macro-F1 score, † denotes the model based on our framework. Average results over 10 runs are reported, best scores for each baseline are in bold.

Model	Sem15		Sem16	
	Acc. (%)	F1 (%)	Acc. (%)	F1 (%)
TD-LSTM	76.39	58.70	83.55	60.26
ISF+TD-LSTM†	77.30	60.54	85.62	63.23
CSF+TD-LSTM†	78.84	61.01	86.32	65.07
GSF+TD-LSTM†	77.66	61.96	86.97	65.77
ATAE-LSTM	78.48	59.77	84.19	62.89
ISF+ATAE-LSTM†	79.80	62.01	85.83	64.37
CSF+ATAE-LSTM†	81.42	63.49	86.34	66.20
GSF+ATAE-LSTM†	80.67	63.85	87.53	65.46
MemNet	77.59	57.47	84.35	61.71
ISF+MemNet†	79.68	60.70	86.74	63.68
CSF+MemNet†	80.52	63.16	87.53	65.62
GSF+MemNet†	81.30	63.79	88.23	66.52
RAM	79.17	62.29	83.88	63.73
ISF+RAM†	81.33	63.64	86.58	65.25
CSF+RAM†	82.37	64.70	87.20	67.37
GSF+RAM†	82.82	65.87	88.72	68.35
IAN	76.84	58.55	84.82	60.65
ISF+IAN†	79.71	61.91	86.26	62.25
CSF+IAN†	81.38	62.90	86.42	64.16
GSF+IAN†	80.81	63.12	88.02	63.96
AOA	77.38	58.79	85.94	63.05
ISF+AOA†	80.72	61.26	87.42	64.33
CSF+AOA†	82.63	63.25	88.54	66.87
GSF+AOA†	82.35	63.49	88.90	67.12
MGAN	78.26	58.46	85.14	59.87
ISF+MGAN†	80.24	60.62	87.46	64.63
CSF+MGAN†	82.05	62.08	87.53	67.61
GSF+MGAN†	83.13	62.76	88.50	68.62
GCAE	77.55	57.43	84.66	60.89
ISF+GCAE†	79.17	60.70	85.94	64.37
CSF+GCAE†	80.67	61.27	87.20	65.94
GSF+GCAE†	81.78	61.93	87.86	66.35
AEN	78.02	60.12	85.78	64.35
ISF+AEN†	79.87	61.99	86.90	65.38
CSF+AEN†	80.92	63.54	88.13	67.20
GSF+AEN†	82.10	64.29	87.64	68.12
BERT	82.41	64.35	88.60	73.62
ISF+BERT†	84.12	65.87	89.25	74.78
CSF+BERT†	84.62	66.53	91.29	75.57
GSF+BERT†	85.04	67.98	91.73	76.12

with existing ABSA models and achieve the state-of-the-art performance for predicting aspect-based sentiment. One main reason is that the proposed framework can leverage aspect-invariant sentiment features from cross-aspect samples for better learning sentiment features of those aspects with limited annotated samples.

In this paper, we explore three schemes of the proposed framework to demonstrate the versatility of our method. Experimental results show that all the three structures can improve the performance of aspect-based sentiment prediction over baseline models. It can also be observed that the concatenation structure (**CSF**) and the gated fusion structure (**GSF**) perform better than the integration structure (**ISF**) on both Sem15 and Sem16 datasets. This indicates that employing an independent feature extractor \mathcal{F} for the extraction of aspect-invariant features is more effective compared to simply using an existing ABSA model. Overall, the gated fusion structure (**GSF**) performs better than other two schemes, showing that aspect-invariant sentiment features are better captured by the gating mechanism.

4.4 Ablation Study

To investigate the impact of different components of our proposed adversarial multi-task learning framework, we conduct experiments based on RAM model and report the results of different structures⁶. As shown in Table 2, Simply incorporation fake samples (w/ fake) only leads to marginal improvements since some of the fake samples may introduce noise to the training process. When integrating multi-task learning (w/ multi) into our framework, the performance can be preeminently improved for all three submodules (w/ a , w/ s and w/ $a+s$). Among them, models with fake sample discrimination (w/ s) achieve more noticeable improvement in comparison with those without (w/ a), which demonstrates that incorporating fake sample discrimination can reduce the negative impact of noisy fake samples and better extract aspect-invariant sentiment features. In addition, compared with models with single-task discrimination, the multi-task learning models with both aspect discrimination and fake sample discrimination achieve better results in general for all the three structures with more significant results for both **CSF** and **GSF**, which demonstrates the effectiveness of the proposed multi-task learning for ABSA. Here, the reason why significant improvement can not be achieved by **ISF** with multi-task learning may be that, it is laborious for a single model to learn various discrepant features concurrently.

It is also worth noticing that significant improvement can be achieved by multi-task learning integrated with adversarial training (w/ adv.+multi+fake w/ $a+s$). This verifies that compared with pure multi-task learning, adversarial multi-task learning can better disentangle aspect-dependent sentiment features and noisy features from aspect-invariant sentiment features by adversarial training, which eventually leads to superior performance in ABSA.

4.5 Impact of the Training Data Size

To further demonstrate that our adversarial multi-task learning framework can extract aspect-invariant sentiment features to improve

Table 2: Experimental results of different variants based on ISF, CSF and GSF structures. a and s represent aspect discrimination and fake sample discrimination respectively. “fake”, “multi” and “adv.” represent fake samples, multi-task learning and adversarial training respectively. * denotes the complete framework.

Model			Sem15		Sem16	
			Acc. (%)	F1 (%)	Acc. (%)	F1 (%)
baseline			79.17	62.29	83.88	63.73
w/ fake			79.30	62.54	84.03	63.85
ISF	w/ multi	w/ s	80.21	63.02	85.97	64.45
	+fake	w/ a	79.86	62.73	85.62	64.23
		w/ $a+s$	80.45	63.08	85.83	64.37
	w/ adv.	w/ s	81.12	63.85	86.62	64.97
	+multi	w/ a	80.97	63.50	86.53	64.73
	+fake	w/ $a+s$ *	81.33	63.64	86.58	65.25
CSF	w/ multi	w/ s	80.03	63.12	86.40	64.78
	+fake	w/ a	79.78	62.88	85.95	64.02
		w/ $a+s$	80.12	63.05	86.57	65.10
	w/ adv.	w/ s	81.27	63.39	87.06	66.73
	+multi	w/ a	80.83	63.10	86.85	65.82
	+fake	w/ $a+s$ *	82.37	64.70	87.20	67.37
GSF	w/ multi	w/ s	80.15	63.27	86.43	65.02
	+fake	w/ a	79.63	62.95	85.82	64.13
		w/ $a+s$	80.23	63.64	86.82	65.35
	w/ adv.	w/ s	81.32	64.68	88.15	67.96
	+multi	w/ a	80.67	64.07	87.56	66.87
	+fake	w/ $a+s$ *	82.82	65.87	88.72	68.35

aspect-based sentiment prediction for aspects with limited annotated samples, we conduct experiments by employing different proportions of annotated data from training datasets and incorporating different proportions of fake cross-aspect samples. The results are shown in Figure 4. According to Figure 4(a) and (b), we can observe that all the three proposed structures of our framework (**ISF**, **CSF** and **GSF**) achieve better performance in comparison with the baseline on both Sem15 and Sem16 dataset. For small proportion of annotated data ($< 40\%$), our three structures can still achieve remarkable performance (**CSF** and **GSF** in particular), that is, when the annotated data is critically insufficient, our framework can still achieve appreciable performance. This implies that extracting aspect-invariant sentiment features can significantly improve the performance of predicting aspect-based sentiment, especially for those aspects with limited annotated data.

In addition, Figure 4(c) and (d) show the performance of using different proportions of fake cross-aspect samples (here, the total number of synthetic cross-aspect samples is 4,873 on Sem15 and 6,310 on Sem16). We observe that the performance of the baseline (the original ABSA model) fluctuates with the increasing number of fake cross-aspect samples. This shows that simply adding fake training samples may introduce noise to the model as some of the

⁶To demonstrate the generalizability of the proposed framework, from here on, we adopt RAM as ABSA model \mathcal{M} in the following experiments, while similar results can also be achieved by regarding other baselines (including BERT) as \mathcal{M} .

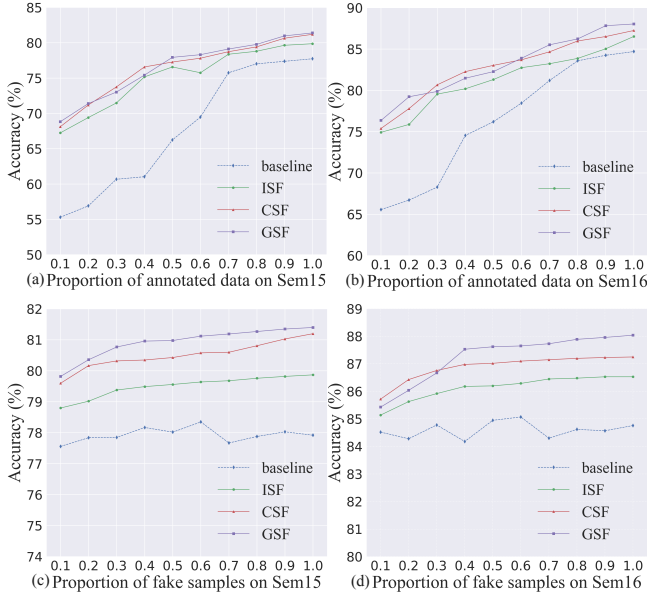
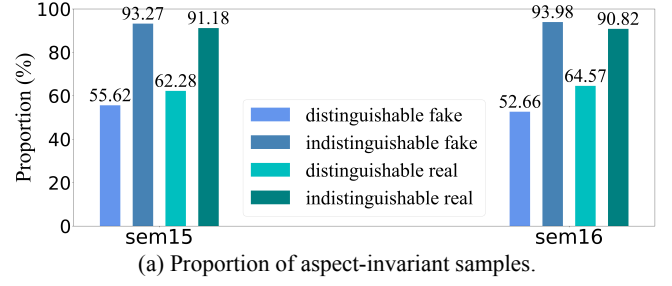


Figure 4: Performance of using different proportions of annotated data on Sem15 and Sem16. (a) and (b) show different percentages of annotated data. (c) and (d) show different percentages of fake cross-aspect samples.

generated synthetic training instances may contain inconsistent sentiment features, which are vulnerable to slash the learning ability of the model. Our three structures, on the contrary, achieve best results with the increase number of fake samples. The performance improvement is more noticeable when the size of fake samples is small ($< 40\%$). This indicates that the proposed adversarial multi-task learning framework can extract aspect-invariant sentiment features from fake cross-aspect samples more effectively, and essentially filtering out noisy signals automatically.

4.6 Detailed Results for Aspect-Invariant Sentiment Extraction

In this subsection, we conduct detailed experiments to verify that the fake cross-aspect samples can be fed into the proposed framework to learn aspect-invariant sentiment features for improving aspect-based sentiment prediction, particularly for aspects with limited annotated data. More concretely, we manually annotate sentiment expressions (i.e. aspect-invariant or aspect-dependent) of samples identified by the multi-task discriminator \mathcal{D} . As demonstrated in Figure 5(a), among cross-aspect samples which are identified wrongly by the discriminator \mathcal{D} (indistinguishable fake/real), over 90% of them contain aspect-invariant sentiment expressions on both Sem15 and Sem16. Examples of such cross-aspect training samples are shown in the first row of Figure 5(b). On the contrary, among those cross-aspect samples which are identified correctly by our discriminator (distinguishable fake/real), only about half of them contain aspect-invariant sentiment expressions. Some aspect-dependent fake examples are shown in the second row of Figure 5(b). Clearly, these fake samples need to be filtered as otherwise, they will confuse the learning of ABSA model.



(a) Proportion of aspect-invariant samples.

Type	Real sample	Fake sample	Result
aspect-invariant	I love this place	I love this ambience I love this food	real ✓
aspect-dependent	service was slow	ambience was slow restaurant was slow	fake ✗

(b) Discrimination results of typical fake samples.

Figure 5: Results of aspect-invariant sentiment features extracted by our framework. ✓ represents that samples contain aspect-invariant sentiment expressions, which could be extracted to improve aspect-based sentiment prediction for aspects with limited annotated data. ✗ represents that the generated samples contain aspect-dependent sentiment expressions, which introduce noise and should be filtered by our framework.

4.7 Visualizations and Qualitative Analysis

To qualitatively demonstrate how the proposed framework improves the performance of ABSA, we visualize the intermediate vectors extracted by the feature extractor \mathcal{F} via t-SNE [26] and analyze what sentiment features are learned from the real testing instances which are indistinguishable by the discriminator. The results are reported in Figure 6. We can observe from Figure 6(a) and (b) that the intermediate vectors from indistinguishable and distinguishable samples are clearly separated by the discriminator \mathcal{D} . In addition, Figure 6(c) and (d) demonstrate that the proportion of indistinguishable instances containing aspect-invariant sentiment expressions is near 90% in both Sem15 and Sem16, and about 90% of the aspect-dependent sentiment features are discriminated by the discriminator \mathcal{D} on both Sem15 and Sem16. This verifies that the proposed framework can effectively distinguish and extract aspect-invariant sentiment features by adversarial multi-task learning.

4.8 Analysis of Gated Fusion Structure (GSF)

As reported earlier, the Gated Fusion Structure (GSF) achieves the best performance in all comparison experiments. To further analyze how the fusion gate better incorporates aspect-invariant sentiment features for improving aspect-based sentiment prediction, we demonstrate in Figure 7 the visualization of intermediate vectors learned by the fusion gate and the distribution of vector values. From Figure 7(a) and (b), we can observe that aspect-invariant and aspect-dependent sentiment features are better separated for both Sem15 and Sem16, which verifies the effectiveness of using the proposed GSF scheme for better incorporating aspect-invariant sentiment features. In addition, Figure 7(c) and (d) show that the peak values

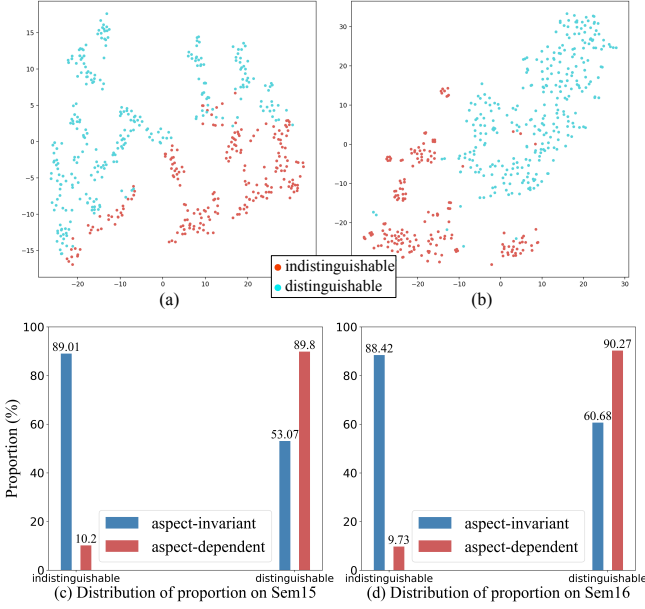


Figure 6: Visualizations of intermediate vectors. (a) and (b) demonstrate intermediate vectors output by the feature extractor \mathcal{F} on Sem15 and Sem16 respectively. Red dots represent those ‘indistinguishable’ real samples that can not be discriminated by a well-trained discriminator \mathcal{D} , cyan dots represent ‘distinguishable’ real samples by the discriminator. (c) and (d) show the distribution of real testing instances containing aspect-invariant and aspect-dependent sentiment expressions on Sem15 and Sem16 respectively. Blue bars denote the proportion of aspect-invariant instances in ‘indistinguishable’ and ‘distinguishable’ categories. Red bars denote the proportion of aspect-dependent instances of these two categories in all aspect-dependent instances.

of intermediate vectors of indistinguishable samples are larger than distinguishable samples on both Sem15 and Sem16. And further, most values of the distribution from indistinguishable samples fall to the right of the peak value, while the values of the distinguishable gather at the left of the peak value. This verifies that the fusion gate of the **GSF** framework can boost the transfer of aspect-invariant sentiment by increasing the value of elements of the intermediate vectors.

4.9 Effect of Off-the-shelf Aspect-Invariant Data

To explore the availability of the aspect-invariant sentiment features for improving the prediction of aspect-based sentiment in the task of ABSA, we output the aspect-invariant samples learned by our feature extractor \mathcal{F} and feed them with different proportions into all baseline models. Here, we adopt different baseline models as \mathcal{M} to train the framework, and select the top 300 aspect-invariant samples for the Sem15 dataset and 400 aspect-invariant samples for the Sem16 dataset from the generated aspect-invariant data. The experimental results are demonstrated in Figure 8. We can observe that with the increasing number of aspect-invariant samples, the

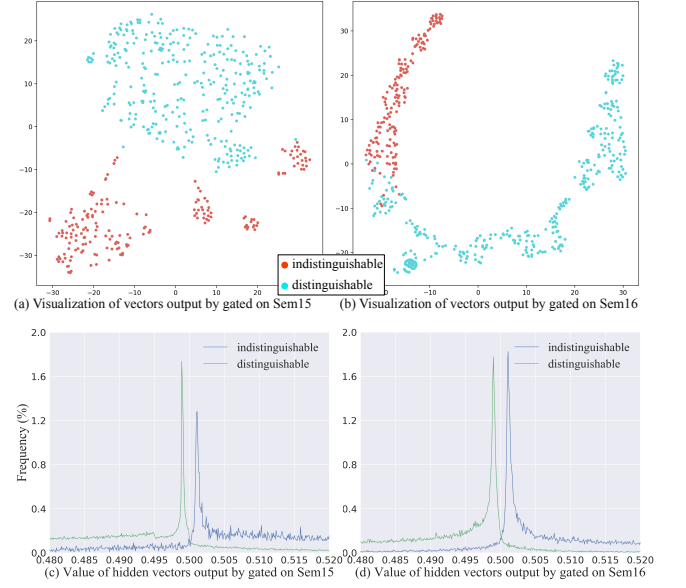


Figure 7: Visualizations and value distributions of intermediate vectors learned by the fusion gate.




Figure 8: Performance of different proportions of aspect-invariant samples on Sem15 and Sem16.

accuracy of aspect-based sentiment prediction increases for all of the baseline models on both Sem15 and Sem16 datasets. This indicates that the proposed framework can indeed extract aspect-invariant samples, and feed them as additional annotated data into existing ABSA models for the improvement of ABSA performance. One possible reason is that aspect-invariant sentiment features would benefit the learning of a better sentiment classifier for those aspects with limited annotated data.

5 CONCLUSION

In this paper, we have proposed a novel adversarial multi-task learning framework to extract aspect-invariant sentiment features from cross-aspect data via adversarial training. The generated fake training instances containing aspect-invariant sentiment features can effectively boost ABSA performance, especially for aspects with limited annotated data. Experimental results on two benchmark datasets

show that the proposed framework can be easily combined with existing neural network-based ABSA models and capture aspect-invariant sentiment features effectively for improving ABSA performance without requiring additional annotated data, and thereby achieve state-of-the-art performance in the ABSA task.

REFERENCES

- [1] Peng Chen, Zhongqian Sun, Lidong Bing, and Wei Yang. 2017. Recurrent Attention Network on Memory for Aspect Sentiment Analysis. In *EMNLP*. 452–461.
- [2] Xilun Chen, Yu Sun, Ben Athiwaratkun, Claire Cardie, and Kilian Weinberger. 2018. Adversarial Deep Averaging Networks for Cross-Lingual Sentiment Classification. *TACL* 6 (2018), 557–570.
- [3] Zhuang Chen and Tiejun Qian. 2019. Transfer Capsule Network for Aspect Level Sentiment Classification. In *ACL*. 547–556.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*. 4171–4186.
- [5] Li Dong, Furu Wei, Chuanki Tan, Duyu Tang, Ming Zhou, and Ke Xu. 2014. Adaptive Recursive Neural Network for Target-dependent Twitter Sentiment Classification. In *ACL*. 49–54.
- [6] Chunling 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-IJCNLP*. 5488–5497.
- [7] Feifan Fan, Yansong Feng, and Dongyan Zhao. 2018. Multi-grained Attention Network for Aspect-Level Sentiment Classification. In *EMNLP*. 3433–3442.
- [8] Devamanyu Hazarika, Soujanya Poria, Prateek Vij, Gangeshwar Krishnamurthy, Erik Cambria, and Roger Zimmermann. 2018. Modeling inter-aspect dependencies for aspect-based sentiment analysis. In *NAACL*. 266–270.
- [9] Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2018. Exploiting Document Knowledge for Aspect-level Sentiment Classification. In *ACL*. 579–585.
- [10] 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*. 504–515.
- [11] Xiaochen Hou, Jing Huang, Guangtao Wang, Kevin Huang, Xiaodong He, and Bowen Zhou. 2019. Selective Attention Based Graph Convolutional Networks for Aspect-Level Sentiment Classification. *arXiv preprint arXiv:1910.10857* (2019).
- [12] Binxuan Huang and Kathleen Carley. 2019. Syntax-Aware Aspect Level Sentiment Classification with Graph Attention Networks. In *EMNLP-IJCNLP*. 5469–5477.
- [13] Binxuan Huang, Yanglan Ou, and Kathleen M Carley. 2018. Aspect level sentiment classification with attention-over-attention neural networks. In *SBP-BRIMS*. 197–206.
- [14] 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-IJCNLP*. 6281–6286.
- [15] Akbar Karimi, Leonardo Rossi, Andrea Prati, and Katharina Full. 2020. Adversarial training for aspect-based sentiment analysis with BERT. *arXiv preprint arXiv:2001.11316* (2020).
- [16] Xin Li, Lidong Bing, Wai Lam, and Bei Shi. 2018. Transformation Networks for Target-Oriented Sentiment Classification. In *ACL*. 946–956.
- [17] Xinlong Li, Xingyu Fu, Guanglu Xu, Yang Yang, Jiuniu Wang, Li Jin, Qing Liu, and Tianyuan Xiang. 2020. Enhancing BERT Representation With Context-Aware Embedding for Aspect-Based Sentiment Analysis. *IEEE Access* 8 (2020), 46868–46876.
- [18] Zheng Li, Xin Li, Ying Wei, Lidong Bing, Yu Zhang, and Qiang Yang. 2019. Transferable End-to-End Aspect-based Sentiment Analysis with Selective Adversarial Learning. In *EMNLP-IJCNLP*. 4582–4592.
- [19] Zheng Li, Ying Wei, Yu Zhang, Xiang Zhang, and Xin Li. 2019. Exploiting coarse-to-fine task transfer for aspect-level sentiment classification. In *AAAI*, Vol. 33. 4253–4260.
- [20] Bin Liang, Jiachen Du, Ruifeng Xu, Binyang Li, and Hejiao Huang. 2019. Context-aware Embedding for Targeted Aspect-based Sentiment Analysis. In *ACL*. 4678–4683.
- [21] Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies* 5, 1 (2012), 1–167.
- [22] Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2017. Adversarial Multi-task Learning for Text Classification. In *ACL*. 1–10.
- [23] Qiao Liu, Haibin Zhang, Yifu Zeng, Ziqi Huang, and Zufeng Wu. 2018. Content attention model for aspect based sentiment analysis. In *WWW*. 1023–1032.
- [24] Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. Multi-Task Deep Neural Networks for Natural Language Understanding. In *ACL*. 4487–4496.
- [25] Dehong Ma, Sujian Li, Xiaodong Zhang, and Houfeng Wang. 2017. Interactive Attention Networks for Aspect-Level Sentiment Classification. In *IJCAI*. 4068–4074.
- [26] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *JMLR* 9, Nov (2008), 2579–2605.
- [27] Navonil Majumder, Soujanya Poria, Alexander Gelbukh, Md Shad Akhtar, Erik Cambria, and Asif Ekbal. 2018. IARM: Inter-aspect relation modeling with memory networks in aspect-based sentiment analysis. In *EMNLP*. 3402–3411.
- [28] Thien Hai Nguyen and Kiyoaki Shirai. 2015. PhraseRNN: Phrase Recursive Neural Network for Aspect-based Sentiment Analysis. In *EMNLP*. 2509–2514.
- [29] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global Vectors for Word Representation. In *EMNLP*. 1532–1543.
- [30] Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphee De Clercq, Veronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Núria Bel, Salud María Jiménez-Zafra, and Gülşen Eryigit. 2016. SemEval-2016 Task 5: Aspect Based Sentiment Analysis. In *SemEval*. 19–30.
- [31] Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. SemEval-2015 Task 12: Aspect Based Sentiment Analysis. In *SemEval*. 486–495.
- [32] Sebastian Ruder, Parsa Ghaffari, and John G. Breslin. 2016. A Hierarchical Model of Reviews for Aspect-based Sentiment Analysis. In *EMNLP*. 999–1005.
- [33] Tao Shen, Xiubo Geng, Tao QIN, Daya Guo, Duyu Tang, Nan Duan, Guodong Long, and Daxin Jiang. 2019. Multi-Task Learning for Conversational Question Answering over a Large-Scale Knowledge Base. In *EMNLP-IJCNLP*. 2442–2451.
- [34] Youwei Song, Jiahai Wang, Tao Jiang, Zhiyue Liu, and Yanguai Rao. 2019. Attentional encoder network for targeted sentiment classification. *arXiv preprint arXiv:1902.09314* (2019).
- [35] Chi Sun, Luyao Huang, and Xipeng Qiu. 2019. Utilizing BERT for Aspect-Based Sentiment Analysis via Constructing Auxiliary Sentence. In *NAACL*. 380–385.
- [36] Kai Sun, Richong Zhang, Samuel Mensah, Yongyi Mao, and Xudong Liu. 2019. Aspect-Level Sentiment Analysis Via Convolution over Dependency Tree. In *EMNLP-IJCNLP*. 5679–5688.
- [37] Duyu Tang, Bing Qin, Xiaocheng Feng, and Ting Liu. 2016. Effective LSTMs for Target-Dependent Sentiment Classification. In *COLING*. 3298–3307.
- [38] Duyu Tang, Bing Qin, and Ting Liu. 2016. Aspect Level Sentiment Classification with Deep Memory Network. In *EMNLP*. 214–224.
- [39] Weichao Wang, Shi Feng, Wei Gao, Daling Wang, and Yifei Zhang. 2018. Personalized Microblog Sentiment Classification via Adversarial Cross-lingual Multi-task Learning. In *EMNLP*. 338–348.
- [40] Wenya Wang, Sinno Jialin Pan, Daniel Dahlmeier, and Xiaokui Xiao. 2016. Recursive Neural Conditional Random Fields for Aspect-based Sentiment Analysis. In *EMNLP*. 616–626.
- [41] Yequan Wang, Minlie Huang, Xiaoyan Zhu, and Li Zhao. 2016. Attention-based LSTM for Aspect-level Sentiment Classification. In *EMNLP*. 606–615.
- [42] Lianwei Wu, Yuan Rao, Haolin Jin, Ambreen Nazir, and Ling Sun. 2019. Different Absorption from the Same Sharing: Sifted Multi-task Learning for Fake News Detection. In *EMNLP-IJCNLP*. 4636–4645.
- [43] Yi Wu, David Bamman, and Stuart Russell. 2017. Adversarial Training for Relation Extraction. In *EMNLP*. 1778–1783.
- [44] Wei Xue and Tao Li. 2018. Aspect Based Sentiment Analysis with Gated Convolutional Networks. In *ACL*. 2514–2523.
- [45] Chao Yang, Hefeng Zhang, Bin Jiang, and Keqin Li. 2019. Aspect-based sentiment analysis with alternating coattention networks. *IPM* 56, 3 (2019), 463–478.
- [46] Chen Zhang, Qiuchi Li, and Dawei Song. 2019. Aspect-based Sentiment Classification with Aspect-specific Graph Convolutional Networks. In *EMNLP-IJCNLP*. 4567–4577.