Cross-Domain Sentiment Classification using Semantic Representation

Shichen Li, Zhongqing Wang*, Xiaotong Jiang and Guodong Zhou
Natural Language Processing Lab, Soochow University, Suzhou, China
{scli_21,devjiang}@outlook.com
{wangzq,gdzhou}@suda.edu.cn

Abstract

Previous studies on cross-domain sentiment classification depend on the pivot features or utilize the target data for representation learning, which ignore the semantic relevance among different domains. To this end, we exploit Abstract Meaning Representation (AMR) to help with cross-domain sentiment classification. Compared with the textual input, AMR reduces data sparsity and explicitly provides core semantic knowledge and correlations among different domains. In particular, we develop an algorithm to construct a sentiment-driven semantic graph from sentence-level AMRs. We further design two strategies to linearize the semantic graph and propose a text-graph interaction model to fully explore the correlations between the text and semantic graph representations for cross-domain sentiment classification. Empirical studies show the effectiveness of our proposed model over several strong baselines. The results also indicate the importance of the proposed sentiment-driven semantic graph for cross-domain sentiment classification.

1 Introduction

As an essential task in natural language processing, sentiment classification has gained great attention from both industry and academia with its wild-spread applications. While, a large number of review domains (e.g., product categories in Amazon) make it intractable to manually annotate enough data in each domain for training domain-specific models. Thus developing automatically cross-domain methods is imperative in this area.

Recent efforts on cross-domain sentiment classification can be separated into three categories: features-based approaches, discriminator-based approaches, and pretrained-based approaches. Feature-based approaches (Blitzer et al., 2007; Yu and Jiang, 2016; Ziser and Reichart, 2019) utilize a key intuition that domain-specific features could be aligned with the help of domain invariant features. Meanwhile, discriminator-based approaches (He et al., 2018; Du et al., 2020; Xue et al., 2020) aim to determine the diversity between domains and predict the polarities of instances holistically. More recently, pre-trained language models are proposed to learn the cross-domain representations with mixed classification and masked language models (Zhou et al., 2020; Karouzos et al., 2021; Wu and Shi, 2022).

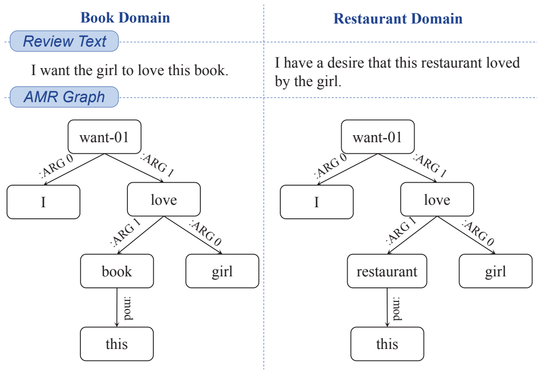

Figure 1: Example of review texts with AMR graphs in different domains.

Despite giving strong empirical results, previous studies ignore the semantic relevance between different domains. As shown in Figure 1, the expression of reviews from different domains are quite different, but they have the similar Abstract Meaning Representation (AMR) (Banarescu et al., 2013) graphs. To this end, we propose a novel framework with AMR for cross-domain sentiment classification. AMR-based semantic graph models the review sentences using a rooted directed acyclic graph, which highlights its main concepts and semantic relations while abstracting away function words (Xu et al., 2020; Bevilacqua et al., 2021; Lyu et al., 2021). It can thus potentially offer core concepts and explicit structures

* Corresponding author

Findings of the Association for Computational Linguistics: EMNLP 2022, pages 289–299

needed for aggregating the meaning of texts in different domains.

In addition, we propose a sentiment-driven graph construction algorithm to better handle long review text and to make better use of the sentimental information. Especially, the process of the construction algorithm can be separated into two stages: document-level semantic graph construction from the sentence-level AMR graphs, and enriching the document-level semantic graph with sentimental information.

We further design two strategies to linearize the sentiment-driven semantic graph: DFS-based linearization is closely related to the way natural language syntactic trees are linearized, and BFS-based linearization enforces a locality principle by which things belonging together are close to each other in the flat representation. Afterward, we propose a text-graph interaction model to fuse the text and the linearized sentiment-driven semantic graph representations for cross-domain sentiment classification. The detailed evaluation shows that our model significantly advances the performance on several benchmark datasets. The results also show that the proposed sentiment-driven semantic graph is beneficial for cross-domain sentiment classification.

The major contributions of this paper are summarized as follows,

- We construct an AMR-based sentiment-driven semantic graph for cross-domain sentiment classification, which makes it easier to find the semantic and sentimental contents in different domains.
- We design two strategies to linearize the sentiment-driven semantic graph, and propose a text-graph interaction model to fully explore the correlations between text and the semantic graph.
- Comprehensive empirical studies show the effectiveness of the proposed model, and also indicate the importance of sentiment-driven semantic graph for cross-domain sentiment classification.

2 Related Works

In this part, we introduce two related topics of this study: cross-domain sentiment classification and abstract meaning representation.

2.1 Cross-Domain Sentiment Classification

Cross-domain sentiment classification has been a long standing attractive research topic due to its real applications where labeled data is only available in a source domain. Previous studies can be separated into three categories: features-based approaches, discriminator-based approaches, and pre-training based approaches.

Feature-based approaches utilize a key intuition that domain-specific features could be aligned with the help of domain invariant features. They are always two-stages approaches, first heuristically select domain-shared pivot features and then use them to learn the correspondence of domain-specific sentiment words (Blitzer et al., 2007; Yu and Jiang, 2016; Ziser and Reichart, 2019; Miller, 2019; Ben-David et al., 2020).

Discriminator-based approaches are end-to-end approaches, which determine the diversity between domains, and predict the polarity of instance directly (He et al., 2018). Recently, the studies of discriminator-based approaches employed adversarial learning (Ganin et al., 2016) to train domain-variant feature extractors and sentiment classifiers holistically, without relying on pivot features (Qu et al., 2019; Du et al., 2020; Xue et al., 2020).

More recently, researches focused on employing pretrained language models in cross-domain classification scenario (Zhou et al., 2020). For example, Karouzos et al. (2021) employed mixed classification and masked cross-domain language model loss to fine-tune the pre-trained model. Wu and Shi (2022) adopt separate soft prompts to learn the vectors for different domains with domain adversarial training strategy.

2.2 Abstract Meaning Representation

Recently, Abstract Meaning Representation (Banarescu et al., 2013) has become an influential formalism for capturing the meaning of a given sentence within a semantic graph (van Noord and Bos, 2017; Xu et al., 2020; Bevilacqua et al., 2021; Bai et al., 2022) and, vice versa, producing text from such a graph (Song et al., 2018; Yao et al., 2020; Jin and Gildea, 2022).

Additionally, AMR's flexibility has resulted in promising improvements for many NLP applications. For example, Liu et al. (2018) parsed text to AMR graphs and transformed them into summary graph for text summarization. Bai et al. (2021)

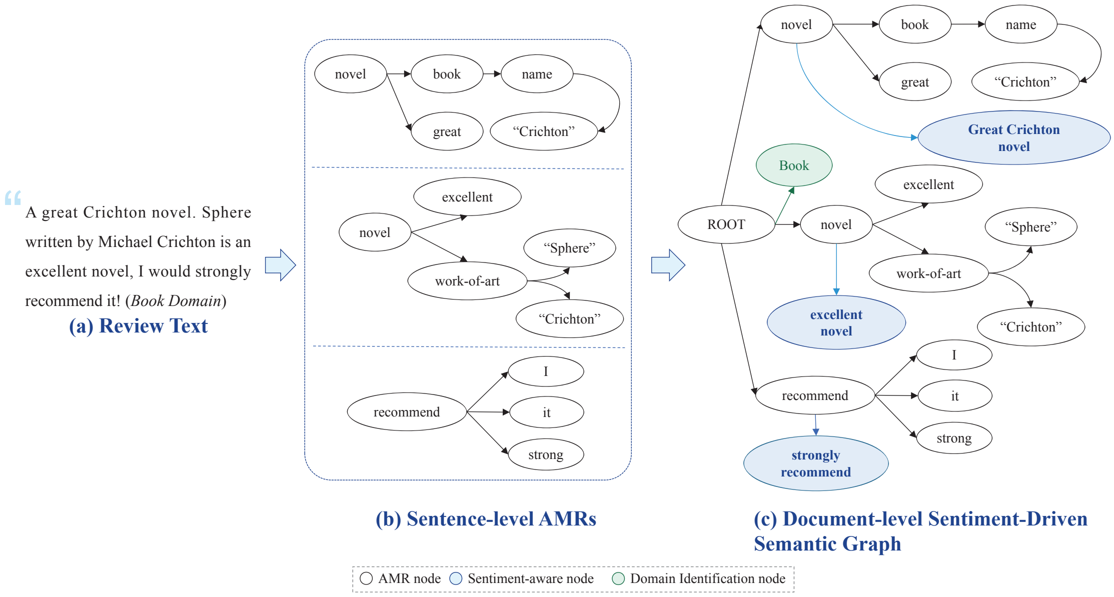

Figure 2: Overview of the sentiment-driven semantic graph construction.

constructed a dialogue-level AMR graph and incorporated it into neural dialogue system for dialog modeling. Zhang et al. (2021) employed AMR to capture semantic structure from complex scientific sentence for biomedical information extraction. Xu et al. (2021) dynamically constructed AMR graph to exploit valid facts for multi-hop science question answering.

Different from previous studies, we employ AMR-based semantic representation for cross-domain sentiment classification. The semantic representation potentially offers core concepts and explicit structures needed for aggregating the meaning of review texts in different domains.

3 Sentiment-Driven Semantic Graph Construction

As shown in Figure 2, we aim to construct a document-level sentiment-driven semantic graph from the sentence-level AMR graphs. Given a review consisting multiple sentences, we apply an off-the-shelf parser[1] of AMR parsing to process these sentences. Afterward, we construct the document-level sentiment-driven semantic graph from these sentence-level semantic graphs. In particular, the process of the document-level sentiment-driven semantic graph construction can be separated into two stages: document-level semantic graph construction, and enriching the semantic graph with sentimental information.

[1] https://github.com/bjascob/amrlib

3.1 Document-level Semantic Graph Construction

The semantic graph of a sentence is represented by a rooted, directed, and acyclic AMR graph (Banarescu et al., 2013). As shown in Figure 2(b), sentence-level AMR graph includes PropBank (Palmer et al., 2005) frames, non-core semantic roles, coreference, entity typing and linking, modality, and negation. The nodes in AMR are concepts instead of words, and the edge types are much more fine-grained compared with traditional semantic languages like dependency parsing and semantic role labeling.

Given a set of sentences and their AMR graphs, we attempt to consolidate all sentence-level graphs to a connected document graph. As shown in Figure 2(c), we first employ a 'ROOT' node to connect the root of each sentence graph, yielding a connected document-level semantic graph.

3.2 Enriching Semantic Graph with Sentimental Information

Since the document-level semantic graph does not contain sentimental information, we extend the original nodes with sentimental information, and let the semantic graph to be a sentiment-driven semantic graph.

As shown in Figure 2(c), we first employ an unsupervised sentiment extraction model (Florescu and Caragea, 2017) to extract sentiment-aware phrases from the review text.

We then attach the sentiment-aware phrases to

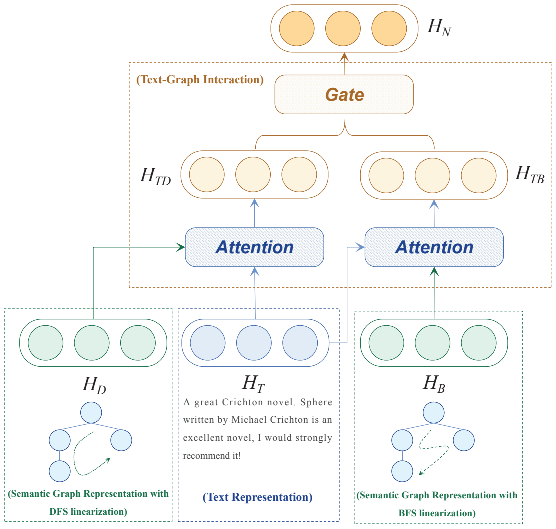

Figure 3: Cross-domain sentiment classification with text-graph interaction model.

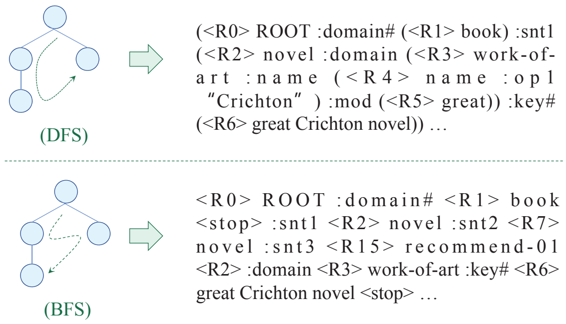

Figure 4: Example of linearzation strategies. The example is based on the sentiment-driven semantic graph in Figure 2(c).

the corresponding nodes. For example, "excellent novel" is attached to node "novel", and "strongly recommend" is attached to node "recommend-01". Meanwhile, we also attach the domain identification node to the ROOT node. For example, the "Book" node in the graph is used to identify that the review text is from book domain.

Therefore, after enriching the document-level semantic graph with sentimental information, we construct a sentiment-driven semantic graph. As shown in Figure 2(c), the new sentiment-driven semantic graph consists of the key concepts and sentiment-driven semantic relations, and it would be very helpful for cross-domain sentiment classification.

4 Cross-Domain Sentiment Classification with Text and Graph Interactions

Our model takes a review text and the corresponding sentiment-driven semantic graph as input. As shown in Figure 3, we first learn the text representation, we then learn the semantic graph representation with different linearization techniques. Thirdly, we employ a text-graph interaction model with self-attention and gate mechanism to learn the fused representation among the text and semantic graph representations. Afterwards, we predict the polarity based on the fused representation.

4.1 Text Representation

We employ BERT (Devlin et al., 2018) to learn the text representation. We first tokenize the review text into a sequence W. Then, we convert each token $w_i \in W$ into vector space by summing the token, segment, and position embeddings. In addition, we use a series of stacked transformer blocks to project the input embeddings into a sequence of contextual vectors. In this way, we obtain the text representation H_T for the review text.

4.2 Semantic Graph Representation

Since it is much easier to integrate a sequence than a graph (Xu et al., 2020), we linearize the sentiment-driven semantic graph to the target sequence. The linearization techniques are fully graph-isomorphic, i.e., it is possible to encode the graph into a sequence of symbols and then decode it back into a graph without losing adjacency information. Following (Bevilacqua et al., 2021), we use the special token "<R0>", "<R1>", ..., "<Rn>" to represent AMR concepts or attributes in the linearized graph and to handle co-referring nodes. Besides, the words start with ":" means different type of edges. The word ":snt1" means "sentence-1", which is a short form of the first sentence. Specially, the words end with "#" like ":domain#" mean the special edge types defined in this paper.

As shown in Figure 4, we employ following two strategies to linearize the sentiment-driven semantic graph:

DFS-based Linearization is quite closely related to the way natural language syntactic trees are linearized (Bevilacqua et al., 2021; Lyu et al., 2021). In particular, we employ a DFS-based traversal algorithm to indicate variables and parentheses to mark visit depth with special tokens. Moreover, we dispose of the redundant slash token (/). These features significantly reduce the length of the output sequence.

BFS-based Linearization enforces a locality principle by which things belonging together are close to each other in the flat representation (Bevilacqua et al., 2021). In addition, BFS is cognitively attractive because it corresponds to a core semantic principle which assumes that the most important pieces of meaning are represented in the upper layers of the graph (Cai and Lam, 2019). To this end, we apply a BFS graph traversal algorithm which starts from the graph root and visits all the children connected by an edge.

Therefore, these two kinds of linearization strategies can capture the semantic and sentimental information from the sentiment-driven semantic graph with either depth or breadth searching. Afterwards, we employ BERT (Devlin et al., 2018) to learn the semantic graph representation H_D and H_B based on these two linearization strategies respectively.

4.3 Text-Graph Interaction

After we learn the text representation H_T and the sentiment-driven semantic graph representations $\{H_D, H_B\}$, we employ a text-graph interaction model to capture the correlations among them. As shown in Figure 3, the model is built based on the self-attention mechanism (Vaswani et al., 2017). It associates source tokens at different positions of the text and the semantic graph representations, by computing the attention score with gate mechanism between each text and semantic graph pair, respectively.

In particular, given the text-graph pairs (H_T, H_D) and (H_T, H_B), we first learn their interaction $\{H_{TD}, H_{TB}\}$ by the text-aware graph attention,

$$H_{TD} = \mathrm{softmax}(\frac{H_D H_T^T}{\sqrt{d_m}})H_T \tag{1}$$

$$H_{TB} = \mathrm{softmax}(\frac{H_B H_T^T}{\sqrt{d_m}})H_T \tag{2}$$

where d_m is the dimension of H_T.

We use a flexible gated mechanism (Li et al., 2018; Zhang et al., 2020) to decide the degree of interactions among text and the semantic graph representations automatically,

$$H_N = \lambda H_{TD} + (1 - \lambda) H_{TB} \tag{3}$$

$\lambda \in (0, 1)$ is calculated by the gated mechanism as below,

$$\lambda = \sigma(W[H_{TD}; H_{TB}] + b) \tag{4}$$

where W and b are learnable model parameters, and $\sigma(\cdot)$ is the sigmoid function. Therefore, we obtain the new representation H_N based on the interactions from the text representation and the sentiment-driven semantic graph representations.

4.4 Cross-domain Sentiment Classification

After we learn the representation H_N based on the interaction among text and semantic graphs, we then employ a multi-layer perceptron model to predict the polarity of it. Formally, given an input vector H_N, a hidden layer is used to induce a set of high-level features as follows:

$$H_P = \sigma(W_p^h H_N + b_p^h), \tag{5}$$

H_P is used as inputs to a softmax output layer:

$$P_P = \mathrm{softmax}(W_p H_P + B_P) \tag{6}$$

where, W_p^h, b_p^h, W_p, and B_p are model parameters.

4.5 Model Training

Our training objective of cross-domain sentiment classification is to minimize the cross-entropy loss over a set of training examples $(d_i, y_i)|_{i=1}^N$, with a ℓ_2-regularization term,

$$\mathcal{L} = -\sum_{i=1}^{N}\sum_{j=1}^{K} y_i \log \hat{y}_i + \frac{\lambda}{2}||\theta_y||^2 \tag{7}$$

where y_i is the pre-defined polarity, $\hat{y}_i$ is the predicted label, θ_y is the set of model parameters and λ is a parameter for ℓ_2-regularization.

5 Experiments

In this section, we first introduce the datasets used for evaluation, and the baseline methods employed for comparison. In addition, we report the experimental results conducted from different perspectives, and analyze the effectiveness of proposed model with different factors.

5.1 Data and Setting

We evaluate our model on a new cross-domain dataset BEAR, which consists of four domains: Books(B), Electronics(E), Airlines(A), and Restaurant(R). The dataset is from Amazon product review dataset (Blitzer et al., 2007), Airline review dataset (Ziser and Reichart, 2018) and Yelp restaurant review dataset (Zhang et al., 2015). For each domain, we choose 800 labeled reviews as

Table 1: Comparison with baselines.

	BERT	PERL	CLIM	UDALM	AdSPT	Ours
B → E	89.4	88.1	89.4	90.7	91.4	**91.6**
B → A	80.9	82.2	81.3	80.8	84.1	**87.4**
B → R	88.7	88.8	88.9	91.6	89.9	**92.3**
E → B	88.9	83.4	89.8	89.4	89.6	**91.1**
E → A	86.5	84.5	86.6	86.1	85.8	**88.1**
E → R	89.7	88.6	89.5	91.1	91.3	**92.6**
A → B	80.5	79.2	80.5	84.9	83.8	**87.1**
A → E	87.2	87.1	87.7	88.5	90.2	**90.8**
A → R	87.4	84.6	87.6	88.2	88.0	**90.0**
R → B	88.2	84.1	88.4	87.4	89.6	**91.3**
R → E	88.6	88.0	88.9	90.4	90.1	**92.1**
R → A	82.1	79.3	80.9	81.7	82.8	**83.4**
Average	86.5	84.8	86.6	87.6	88.1	**89.8**

training data in source domain or as testing data in target domain, and the another 200 reviews as validation data. In addition, we also choose 16,000 unlabeled reviews in each domain. The advantage of the new cross-domain dataset is that the similarity between domains is much smaller than previous datasets.

In this study, we use Jaccard similarity score (Ioffe, 2010) to evaluate domain similarity, which is intuitively calculated as the unique word overlap for all words present in two domains. In short, the lower score indicates the lower similarity between the two domains. The average Jaccard similarity score of the proposed new dataset is 0.243, while the score of Amazon dataset is 0.320. Therefore, the new dataset with larger domain diversity would be more helpful to evaluate the performance of cross-domain methods.

We use the $\text{BERT}_{\text{base}}$[2] and fine-tune its parameters. In particular, we employ the unlabeled reviews from all domains to fit the MLM tasks in the pre-training phase. During the fine-tuning phase, we tune the parameters of our cross-domain sentiment classification model by grid searching on the validation dataset. We select the best model by early stopping using the Accuracy results on the validation dataset. Adam (Kingma and Ba, 2015) is adopted with the learning rate 1×10^{-5}, and the batch size is 6. Our experiments are carried out with an NVIDIA 3090 24G GPU.

The experimental results are obtained by averaging three runs with random initialization, where Accuracy are used as the evaluation metrics.

5.2 Main Results

Table 1 shows the results of different systems. We compare the proposed model with various strong baselines,

- **BERT** is a basic model which simply employs BERT (Devlin et al., 2018) for cross-domain sentiment classification.
- **PERL** (Ben-David et al., 2020) extends contextualized word embedding models of representation learning (i.e., BERT) with pivot-based fine-tuning.
- **UDALM** (Karouzos et al., 2021) employs mixed classification and masked language model loss to fine-tune the pre-trained model. It thus can adapt to the target domain distribution in a robust and sample efficient manner.
- **CLIM** (Li et al., 2021) is a self-training method based on contrastive learning with mutual information maximization. It aims to explore the potential of contrastive learning for domain-invariant and task-discriminative features.
- **AdSPT** (Wu and Shi, 2022) adopts separate soft prompts to learn different vectors for different domains, and also uses a novel domain adversarial training strategy to learn domain-invariant representations between different domains.

[2] https://huggingface.co/bert-base-uncased

Table 2: Impact of sentiment-driven semantic graph.

Method	Accuracy
Ours	89.8
-Graph	86.5
-Enrich	89.1
-Sentiment	89.2
-Domain	89.6

From the table, we find that the previous cross-domain methods (e.g., UDALM and AdSPT) outperform the basic sentiment classification model (i.e., BERT). It indicates that the effectiveness of these cross-domain methods with either pre-training language model or adversarial training strategy for cross-domain sentiment classification.

By contrast, our proposed model outperforms the previous cross-domain methods significantly ($p < 0.05$), since the proposed model employs AMR-based semantic representation for cross-domain sentiment classification. This also proves that the sentiment-driven semantic graph is very useful to exploit the semantic correlations between different domains.

5.3 Effect of Sentiment-Driven Semantic Graph

As shown in Table 2, we employ ablation experiments to analyze the impact of the proposed sentiment-driven semantic graph with average accuracy score in all the domain pairs. If we totally remove sentiment-driven semantic graph (-Graph), it degrades the proposed model to simple BERT-based sentiment classification model, and the performance drops to 86.5%. It shows that sentiment-driven semantic graph can enable a model to better capture the semantic representation of review text, and is beneficial to cross-domain sentiment classification.

In addition, if we remove the sentimental information enriching part (-Enrich) of proposed model, and just keep the document-level semantic graph for cross-domain sentiment classification, the performance drops to 89.1%. It shows that the sentimental information enriching part can improve the performance of the basic AMR-based semantic graph. Furthermore, we also find that both sentimental and domain information are beneficial to construct the sentiment-driven semantic graph. If we remove these two components, the performance drops to 89.2% and 89.6%.

Table 3: Influence of different factors in the proposed model.

Method		Accuracy
BERT		86.5
Linearization	BFS	89.1
	DFS	89.3
Correlation	Concat	88.6
	Attention	89.4
Ours		89.8

5.4 Impact of Model Configuration

This subsection analyzes the influence of different factors in the proposed cross-domain sentiment classification model with average accuracy score in all the domain pairs. As shown in Table 3, "BFS" and "DFS" means that we employ BFS or DFS linearization strategy to linearize the sentiment-driven semantic graph respectively. In addition, "Concat" and "Attention" means that we simply use concatenation or attention mechanism to capture the correlation between text and semantic graph representations. "Ours" is the proposed text-graph interaction model which employ both attention and gate mechanism to model the correlations between text and semantic graph representations.

From the results, we find that both BFS and DFS linearization strategy are beneficial to linearize the AMR-based semantic representation. In addition, we also find that attention mechanism is more effective than simply concatenation method. Furthermore, the proposed text-graph interaction model (Ours) gives the best performance. It indicates that the text-graph interaction model with attention and gate mechanism is more useful to capture the correlations between text and the sentiment-driven semantic graph.

6 Analysis and Discussion

In this section, we give analysis and discussion to show the importance of sentiment-driven semantic graph for cross-domain sentiment classification.

6.1 Influence of Training Data Size

Since the cross-domain sentiment classification task is very sensitive with the training data in source domain, we analyze the size of training data in Figure 5. From the figure, we find that the more training data, the higher performance our proposed model can reach with the average accu-

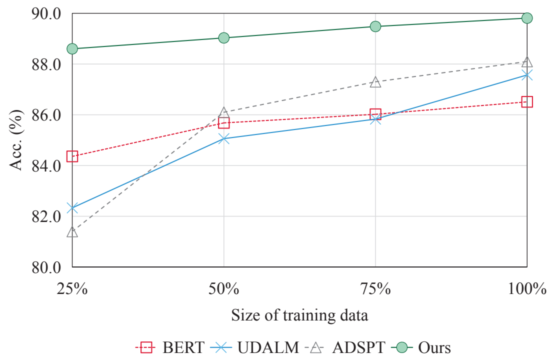

Figure 5: Results with different training data size.

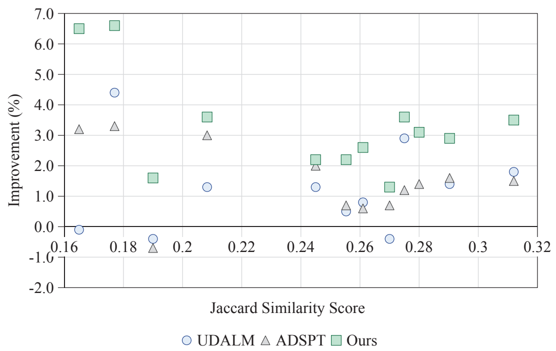

Figure 6: Improvement of different models with Jaccard Score.

Review Text: Harry Potter, it is greatly shocking. I love this extraordinary series. I have every one of the books in this series and will never let them go.
Sentiment-Driven Semantic Graph: 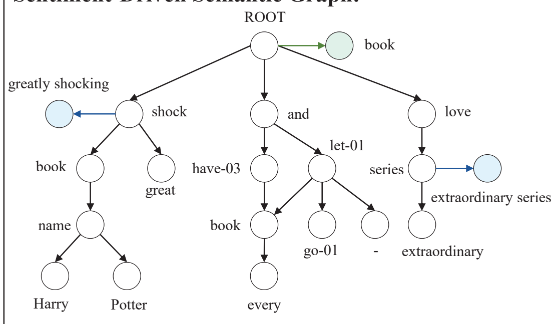
Source Domain: Airline **Target Domain:** Book
BERT: Negative
Ours: Positive

Figure 7: Example of case study.

racy in all domains. In addition, different from the other models, our proposed model can reach a high performance even there are only 25% labeled reviews in source domain. It shows that the semantic representation can capture the high level semantic correlations in different domains, and significantly reduces the size of labeled data in source domain.

6.2 Influence of Similarity between Domains

The performances of cross-domain models are always influenced by the similarity (or diversity) between source and target domains. Therefore, we compare the improvements of three cross-domain models (i.e., UDALM, ADSPT and our proposed model) with the basic BERT model based on different Jaccard similarity score (Ioffe, 2010). The score is used to calculate the similarity between source and target domain in the proposed dataset.

As shown in the figure, our proposed model always outperforms other models whenever the similarity score is large or small. Furthermore, our proposed model can still reach a high improvement even when the similarity score between the two domains is much smaller. It indicates that the proposed sentiment-driven semantic graph can capture the core concepts and explicit structures in different domains, and is very effective for cross-domain sentiment classification.

6.3 Case Study

Figure 7 gives an example to illustrate the effectiveness of the proposed model compared with BERT model. The domain of the given example is Book, while the training data is from Airline. The basic BERT model incorrectly predicts the polarity of the example. It might be confused of BERT model about what *them* refers to in the last clause "*will never let them go*". However, the proposed model successfully predicts the polarity. In the review text, *book* co-referred with *them* is located far away from *let*, but in the AMR graph, *book* is directly linked with *let*. Besides, the sentiment-driven semantic graph enable the sentimental-aware nodes to get close to their target concept nodes, which is much easier for the model to identify the polarity with the guidance of ARM graph. In addition, the result also indicates that the sentiment-driven semantic graph capture both sentimental and semantic information in the review text.

7 Conclusion

In this study, we employ AMR-based semantic representation for cross-domain sentiment classification. AMR can reduce data sparsity, and explicitly provides core semantic knowledge and correlations among different domains. In particular, we develop an algorithm to construct a sentiment-driven semantic graph from sentence-level AMRs. We further design two strategies to linearize the semantic graph and propose a text-graph interaction model to fuse the text and semantic graph representations for cross-domain sentiment classification. Experimental results show the superiority of using the sentiment-driven semantic representations on cross-domain sentiment classification.

Limitations

Although AMR graph has been proved effect for cross-domain sentiment classification, it is still necessary for us to explore a more suitable way to integrate AMR graph for sentiment classification. In addition, the proposed model needs large GPU resources since it should learn both text and semantic representations.

Acknowledgement

We would like to thank Dr. Junhui Li, Yanzhi Xu and Yueying Hua for their helpful advice and discussion during this work. Also, we would like to thank the anonymous reviewers for their insightful and valuable comments. This work is supported by the National Natural Science Foundation of China (No. 62076175, No. 61976146).

References

Xuefeng Bai, Yulong Chen, Linfeng Song, and Yue Zhang. 2021. Semantic representation for dialogue modeling. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 4430–4445. Association for Computational Linguistics.

Xuefeng Bai, Yulong Chen, and Yue Zhang. 2022. Graph pre-training for AMR parsing and generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 6001–6015.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *ACL 2013*, pages 178–186.

Eyal Ben-David, Carmel Rabinovitz, and Roi Reichart. 2020. PERL: pivot-based domain adaptation for pre-trained deep contextualized embedding models. *Trans. Assoc. Comput. Linguistics*, 8:504–521.

Michele Bevilacqua, Rexhina Blloshmi, and Roberto Navigli. 2021. One SPRING to rule them both: Symmetric AMR semantic parsing and generation without a complex pipeline. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 12564–12573.

John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL 2007, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, June 23-30, 2007, Prague, Czech Republic*. The Association for Computational Linguistics.

Deng Cai and Wai Lam. 2019. Core semantic first: A top-down approach for AMR parsing. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3797–3807.

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.

Chunning Du, Haifeng Sun, Jingyu Wang, Qi Qi, and Jianxin Liao. 2020. Adversarial and domain-aware BERT for cross-domain sentiment analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4019–4028. Association for Computational Linguistics.

Corina Florescu and Cornelia Caragea. 2017. Positionrank: An unsupervised approach to keyphrase extraction from scholarly documents. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1105–1115.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. 2016. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.*, 17:59:1–59:35.

Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2018. Adaptive semi-supervised learning for cross-domain sentiment classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3467–3476. Association for Computational Linguistics.

Sergey Ioffe. 2010. Improved consistent sampling, weighted minhash and L1 sketching. In *ICDM*

2010, The 10th IEEE International Conference on Data Mining, Sydney, Australia, 14-17 December 2010, pages 246–255.

Lisa Jin and Daniel Gildea. 2022. Rewarding semantic similarity under optimized alignments for amr-to-text generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 710–715.

Constantinos Karouzos, Georgios Paraskevopoulos, and Alexandros Potamianos. 2021. UDALM: unsupervised domain adaptation through language modeling. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 2579–2590.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Changliang Li, Liang Li, and Ji Qi. 2018. A self-attentive model with gate mechanism for spoken language understanding. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3824–3833.

Tian Li, Xiang Chen, Shanghang Zhang, Zhen Dong, and Kurt Keutzer. 2021. Cross-domain sentiment classification with contrastive learning and mutual information maximization. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*, pages 8203–8207.

Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman M. Sadeh, and Noah A. Smith. 2018. Toward abstractive summarization using semantic representations. *CoRR*, abs/1805.10399.

Chunchuan Lyu, Shay B. Cohen, and Ivan Titov. 2021. A differentiable relaxation of graph segmentation and alignment for AMR parsing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 9075–9091.

Timothy A. Miller. 2019. Simplified neural unsupervised domain adaptation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 414–419.

Martha Palmer, Paul R. Kingsbury, and Daniel Gildea. 2005. The proposition bank: An annotated corpus of semantic roles. *Comput. Linguistics*, 31(1):71–106.

Xiaoye Qu, Zhikang Zou, Yu Cheng, Yang Yang, and Pan Zhou. 2019. Adversarial category alignment network for cross-domain sentiment classification. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2496–2508. Association for Computational Linguistics.

Linfeng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. 2018. A graph-to-sequence model for amr-to-text generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1616–1626.

Rik van Noord and Johan Bos. 2017. Neural semantic parsing by character-based translation: Experiments with abstract meaning representations. *CoRR*, abs/1705.09980.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.

Hui Wu and Xiaodong Shi. 2022. Adversarial soft prompt tuning for cross-domain sentiment analysis. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 2438–2447.

Dongqin Xu, Junhui Li, Muhua Zhu, Min Zhang, and Guodong Zhou. 2020. Improving AMR parsing with sequence-to-sequence pre-training. In *EMNLP 2020*, pages 2501–2511.

Weiwen Xu, Huihui Zhang, Deng Cai, and Wai Lam. 2021. Dynamic semantic graph construction and reasoning for explainable multi-hop science question answering. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 1044–1056. Association for Computational Linguistics.

Qianming Xue, Wei Zhang, and Hongyuan Zha. 2020. Improving domain-adapted sentiment classification by deep adversarial mutual learning. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 9362–9369. AAAI Press.

Shaowei Yao, Tianming Wang, and Xiaojun Wan. 2020. Heterogeneous graph transformer for graph-to-sequence learning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7145–7154.

Jianfei Yu and Jing Jiang. 2016. Learning sentence embeddings with auxiliary tasks for cross-domain sentiment classification. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 236–246. The Association for Computational Linguistics.

Weisheng Zhang, Kaisong Song, Yangyang Kang, Zhongqing Wang, Changlong Sun, Xiaozhong Liu, Shoushan Li, Min Zhang, and Luo Si. 2020. Multi-turn dialogue generation in e-commerce platform with the context of historical dialogue. In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, pages 1981–1990.

Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 649–657.

Zixuan Zhang, Nikolaus Nova Parulian, Heng Ji, Ahmed Elsayed, Skatje Myers, and Martha Palmer. 2021. Fine-grained information extraction from biomedical literature based on knowledge-enriched abstract meaning representation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 6261–6270. Association for Computational Linguistics.

Jie Zhou, Junfeng Tian, Rui Wang, Yuanbin Wu, Wenming Xiao, and Liang He. 2020. Sentix: A sentiment-aware pre-trained model for cross-domain sentiment analysis. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 568–579.

Yftah Ziser and Roi Reichart. 2018. Pivot based language modeling for improved neural domain adaptation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1241–1251. Association for Computational Linguistics.

Yftah Ziser and Roi Reichart. 2019. Task refinement learning for improved accuracy and stability of unsupervised domain adaptation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5895–5906. Association for Computational Linguistics.