

ECNU_ICA at SemEval-2022 Task 10: A Simple and Unified Model for Monolingual and Crosslingual Structured Sentiment Analysis

Qi Zhang¹, Jie Zhou^{2*}, Qin Chen¹, Qingchun Bai³, Jun Xiao³, and Liang He¹

¹School of Computer Science and Technology East China Normal University, China

²School of Computer Science, Fudan University, China

³Shanghai Open University, China

qzhang@stu.ecnu.edu.cn, jzhou@ica.stc.sh.cn

{qchen, lhe}@cs.ecnu.edu.cn, qc_bai@foxmail.com, xiaoj@sou.edu.cn

Abstract

In this paper, we focus on the structured sentiment analysis task that is released on SemEval-2022 Task 10. The task aims to extract the structured sentiment information (e.g., holder, target, expression and sentiment polarity) in a text. We propose a simple and unified model for both the monolingual and crosslingual structured sentiment analysis tasks. We translate this task into an event extraction task by regarding the expression as the trigger word and the other elements as the arguments of the event. Particularly, we first extract the expression by judging its start and end indices. Then, to consider the expression, we design a conditional layer normalization algorithm to extract the holder and target based on the extracted expression. Finally, we infer the sentiment polarity based on the extracted structured information. We conduct the experiments on seven datasets in five languages. It attracted 233 submissions in monolingual subtask and crosslingual subtask from 32 teams. Finally, we obtain the top 5 place on crosslingual tasks.

1 Introduction

The identification of sentiment in the text is an important field of study. Users' opinions on products, events, topics, and so on are valuable for both the company and government to improve products or policies. Recently, more and more researchers focus on fine-grained sentiment analysis tasks, such as structured sentiment analysis, which can be formulated into tuple extraction from the context (Wiebe et al., 2005).

We focus on structured sentiment analysis, which is released by SemEval 2022 task 10. Formally, the task aims to extract all of the opinion tuples $O = \{O_i, \dots, O_n\}$ in a text. A tuple (h, t, e, p) can represent the structure of sentiment in context. As shown in Figure 1, "Some other" (h) is a holder who expresses a polarity "positive" (p)

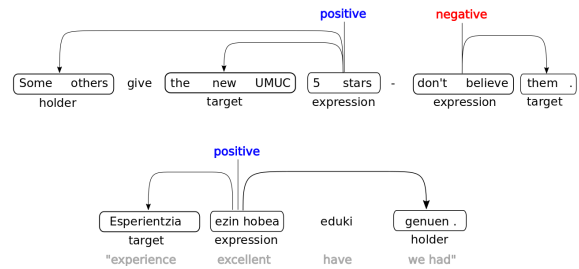


Figure 1: A structured sentiment graph is composed of a holder, target, sentiment expression, their relationships and a polarity attribute. Holders and targets can be null.

towards a target "the new UMUC" (t) through a sentiment expression "5 stars" (e), implicitly defining the relationships between the elements of a sentiment graph. Moreover, the SemEval task 10 is divided into two settings of increasing complexity: **Setting A** trains and predicts the sentiment polarity on monolingual. **Setting B** trains on the source language and tests on the target language. Particularly, we can train on any of the other datasets, as well as any other resource that does not contain sentiment annotations in the target language. The datasets used for this task are provided by the competition organizers, who collected and annotated the corpus (Barnes et al., 2021).

Structured sentiment analysis can be resolved down into four sub-tasks: a) expression extraction, b) target extraction, c) holder extraction, and d) sentiment polarity classification. Previous work on information extraction has used pipeline methods that extract the holders, targets, and expressions (tasks a-c). Most of the works conduct experiments on the MPQA dataset (Wiebe et al., 2005). Choi et al. (2006); Yang and Cardie (2012) used CRFs with the extracted features (e.g., named-entity tagger, sentiment lexicons, and dependency parsers), which results in a very strong baseline. For a complicated task with a small size of the training data, these feature-based mechanism learning techniques often still perform better than neural-based models,

* Jie Zhou is the corresponding author of this paper.

such as Bi-LSTMs (Katiyar and Cardie, 2016). The end-to-end approaches have shown some potential by learning the relationships among the multiple subtasks (Zhang et al., 2019; Zhou et al., 2020a). However, all of these studies ignore the sentiment polarity classification subtask. Barnes et al. (2021) regarded the structured sentiment analysis task as a dependency parsing task to model the relationships among the elements.

In this paper, we transform this task into an event extraction framework. Because we found the expression word is the same as the trigger word in the event extraction, both of them can uniquely represent the whole structured tuple. Besides, we regard the other elements (e.g., holder and targets) as the arguments of the sentiment structured information. Above all, we propose a pipeline model for this task. The key problem for this task is how to model the relationships among the targets, holders and expressions. To take the expressions into account, we introduce a conditional layer normalization method to extract the holders and targets w.r.t. the expression. Our simple and unified model is appropriate for both monolingual and crosslingual tasks. Experiments on datasets of monolingual and crosslingual tasks show our approach is capable to improve performance significantly on structured sentiment analysis as event extraction.

2 Related Work

2.1 Structured Sentiment Analysis

The goal of structured sentiment analysis is to extract the holders, expressions, and polarities w.r.t. the targets (Zhou et al., 2019). Early researches formulated the subtasks of structured sentiment analysis into independent span extraction or relation extraction. Choi et al. (2006) investigated a joint approach to extract entities and relations at the same time for opinion recognition and analysis. Yang and Cardie (2012) proposed semi-Markov conditional random fields (semi-CRFs) to extract opinion expressions at segment level. Katiyar and Cardie (2016) used deep bidirectional LSTMs for joint extraction of opinion entities and relations (e.g., the IS-FROM, IS-ABOUT) that connect the entities to extract expressions with their associated holders and targets. Zhang et al. (2019) presented a transition-based end-to-end method to extract the elements (e.g., holders, targets, and expressions) with their relationships. Then, to take the sentiment polarity into account, systems like IMN (He et al.,

2019), SK-GCN (Zhou et al., 2020b) and RACL (Chen and Qian, 2020) have been developed.

Moreover, Barnes et al. (2021) introduced a parsing-based algorithm that implements a real unified structured sentiment analysis. They regarded structured sentiment analysis as a dependent parsing problem to model the relatedness among the holder, target, expression, and polarity. Different from these works, we translate structured sentiment analysis to an event-extraction framework.

2.2 Multilingual Sentiment Classification

Traditional methods for multilingual sentiment analysis are based on machine translation. Then, neural-based models product the source and target languages into a common space via parallel data or dictionaries. Lample and Conneau (2019) proposed two methods to learn cross-lingual language models: one unsupervised that uses cross language modeling to learn cross language representation. And one supervised that leverages parallel data with a new cross-lingual language model objective. Tellez et al. (2017) trained an SVM classifier by leveraging language-dependent and independent features. However, these machine learning approaches also require a feature extraction phase. We eliminate by incorporating deep learning approaches since they can learn features automatically. Furthermore, (Wan, 2008) designed a new approach to improve Chinese sentiment analysis using reliable English datasets. Recently, multi-lingual pre-trained language models (e.g., mBert (Devlin et al., 2019)) obtains state-of-the-art performance for crosslingual tasks.

3 Our Approach

In this paper, we propose a simple and unified model to extract the structured sentiment analysis. This section describes how we redefine the sentiment classification task to an event extraction task, detail the pipeline method. Fig. 3 shows the framework of our model, which consists of three components, expression extraction, arguments extraction (including holder extraction and target extraction), sentiment classification. First, we extract the expression based on (multi-lingual) pre-trained language models (PLMs) (See Subsection 3.2). Second, we design a conditional layer normalization strategy to extract the holders and targets by incorporating the expression. Finally, we perform sentiment classification tasks based on the sentence




Figure 2: The process of our solution, It consists of three parts: Expression Extraction, Argument Extraction, Sentiment classification.

representation learned by the previous part (See Subsection 3.3 and 3.4).

3.1 Task Definition

Given a sentence $s = \{w_1, \dots, w_{|s|}\}$, where w_i is the i -th word in the sentence s , which contains $|s|$ words. The goal of this task is to extract all the sentiment tuples $O = \{O_1, \dots, O_{|O|}\}$ in the text, where the tuple (e, a, p) consists of expression e , arguments a (target t , holder h) and sentiment polarities c tuple. The event sentiment polarity class $c \in \{P, N, O\}$, which represents positive, negative and neutral.

3.2 Expression Extraction

Through the data analysis, we found that the expression uniquely identifies the tuple O since the targets and holders may not exist in the text. Thus, we extract the expression first by regarding it as the trigger. We first obtain the contextual word embedding based on the PLMs. For the monolingual setting, we use language-specific PLMs. For crosslingual setting, we use the multi-lingual PLMs. Particularly, we input the sentence s into PLMs to obtain the word embedding $X = \{x_1, x_2, \dots, x_{|s|}\}$,

$$X = \text{PLM}(s) \quad (1)$$

Then, we extract the expression by two token-level binary classifiers to predict the start and end indices of the expression. A linear layer with an activation function is adopted as the classifier,

$$\begin{aligned} p_e^s &= \text{Sigmoid}(W_e^s X + b_e^s) \\ p_e^e &= \text{Sigmoid}(W_e^e X + b_e^e) \end{aligned} \quad (2)$$

where p_e^s and p_e^e are the predicted probability distribution of expression e 's start and end indices, $W_e^s, W_e^e, b_e^s, b_e^e$ are the learnable weights.

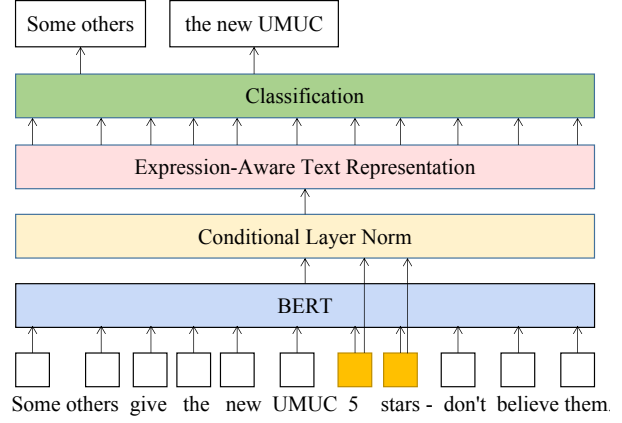


Figure 3: The framework of our model

3.3 Argument Extraction

The key challenge for extracting targets and holders with respect to the extracted expression is how to integrate the expression into the extractor model. Inspired by (De Vries et al., 2017), we design a conditional layer normalization method to take the expression into account. Specifically, we use the representation of the expression to control the β and γ in the layer normalization. In this way, we can obtain the expression-aware representation for target and holder extraction. The word representation h_i after the conditional layer normalization (CLN) is computed as,

$$\begin{aligned} h_i &= \text{CLN}(x_i) = \gamma \frac{e_i - \mu}{\sigma} + \beta, \\ \gamma &= W_\gamma[x_{e^s}; x_{e^e}], \\ \beta &= W_\beta[x_{e^s}; x_{e^e}], \end{aligned} \quad (3)$$

where μ, σ are the mean and standard deviation of the elements in x_i , e^s and e^e are the start and end indices of the expression e , W_γ and W_β are the trainable parameters. We use the concatenation of the expression's start and end word representation as the expression representation.

Based on the expression-aware representation $H = \{h_1, h_2, \dots, h_{|s|}\}$, we train two classifiers to predict the probability distribution of the start and end indices for targets and holders respectively.

$$\begin{aligned} p_t^s &= \text{Sigmoid}(W_t^s H + b_t^s) \\ p_t^e &= \text{Sigmoid}(W_t^e H + b_t^e) \\ p_h^s &= \text{Sigmoid}(W_h^s H + b_h^s) \\ p_h^e &= \text{Sigmoid}(W_h^e H + b_h^e) \end{aligned} \quad (4)$$

where p_t^s/p_h^s and p_t^e/p_h^e are the predicted probability distribution of target t /holder h 's start and end in-

Table 1: The statistics information of datasets.

Dataset	Language	# sents	# holders	# targets	# expr.
NoReC _{Fine}	Norwegian	11437	1128	8923	11115
MultiB _{EU}	Basque	1521	296	1775	2328
MultiB _{CA}	Catalan	1678	235	2336	2756
OpeNer _{es}	Spanish	2057	255	3980	4388
OpeNer _{en}	English	2494	413	3850	4150
MPQA	English	10048	2279	2452	2814
DS _{Unis}	English	2803	86	1119	1119

dices, $W_t^s, W_t^e, b_t^s, b_t^e/W_h^s, W_h^e, b_h^s, b_h^e$ are the learnable weights.

3.4 Sentiment Classification

Finally, we infer the sentiment polarity of the structured tuple (h, t, e) . Since the text representation contains the expression information that can represent the tuple, we use a max pooling operator based on H to obtain the tuple representation r ,

$$r = \text{MaxPooling}(H) \quad (5)$$

We input the tuple representation r into a sentiment classifier to predict the sentiment polarity distribution towards the target,

$$p_c = \text{Softmax}(W_c r + b_c) \quad (6)$$

where W_c and b_c are the parameters.

4 Experiments

In this section, we first give the experimental setup, including datasets, implementation details and evaluation metrics (Section 4.1). Then, we present the experimental results and analysis on both the monolingual and multilingual settings (See Section 4.2).

4.1 Experimental Setup

Datasets To evaluate the effectiveness of our model, we conduct our experiments on datasets of multiple languages, including English, Basque, Catalan, Norwegian. MultiB_{EU}, MultiB_{CA} (Barnes et al., 2018) and NoReC_{Fine} (Øvrelid et al., 2020) are the reviews data in Basque, Catalan, and Norwegian. MPQA (Wiebe et al., 2005) is a English dataset that contains expressions, holders, targets and their relationships. DS_{Unis} (Toprak et al., 2010) contains labeled opinions for user reviews about universities and services. The OpenNER dataset consists of labeled reviews of hotels from the guests. And it’s divided into two languages:

OpeNer_{en} in English, OpeNer_{es} in Spanish. The statistics of the datasets are shown in Table 1. We can find that the size of the labeled data is limited, especially the number of labeled holders. For example, in DS_{Unis} dataset, there are only 86 holders, which limits the performance of the neural models largely.

Implementation Details For the monolingual setting, we utilize the language-specific pre-trained language models (PLMs) as the word embedding, which are downloaded from Hugging Face¹. We finetune the parameters on the training data for this setting. For the crosslingual setting, we use the multi-lingual PLMs (e.g., mBEER, XLM) and fix the parameters on the training phase. We use Adam optimizer with the learning rates of 1e-5. The dimensions of word embedding are 128. The max sequence length is 512. The dropout is 0.1. We train all models for 100 epochs and keep the model that performs best regarding F1 on the dev set. We use default hyperparameters from Kurtz et al. (2020) and run all of our models five times with different random seeds. The reported test results are based on the parameters that obtain the best performance on the development.

Evaluation Metrics Following the previous works (Barnes et al., 2021), as we are interested not only in extraction or classification but rather in the full structured sentiment task, we use Sentiment Graph F1 as the final metric.

4.2 Experimental Results and Analysis

For experimental results, we report the majority baseline for each language. Our unified model can perform on both the monolingual and crosslingual settings and obtain good performance on these tasks, our main experimental results are presented in Table 2 and 3. We follow metrics in (Barnes

¹<https://huggingface.co/models>

Table 2: Top 15 Results for the monolingual setting.

Monolingual							
User	NoReC _{Fine}	MultiB _{CA}	MultiB _{EU}	OpeNer _{en}	OpeNer _{es}	MPQA	DS _{Unis}
zhixiaobao	0.529 (2)	0.728 (1)	0.739 (1)	0.760 (2)	0.722 (4)	0.447 (1)	0.494 (1)
Cong666	0.524 (3)	0.728 (1)	0.739 (1)	0.763 (1)	0.742 (1)	0.416 (2)	0.485 (2)
gmorio	0.533 (1)	0.709 (3)	0.715 (3)	0.756 (3)	0.732 (3)	0.402 (3)	0.463 (3)
colorful	0.504 (4)	0.681 (6)	0.723 (2)	0.747 (4)	0.735 (2)	0.375 (5)	0.410 (9)
whu_stone	0.483 (8)	0.711 (2)	0.681 (6)	0.727 (5)	0.686 (7)	0.379 (4)	0.373 (13)
KE_AI	0.483 (8)	0.711 (2)	0.681 (6)	0.727 (5)	0.686 (7)	0.364 (7)	0.373 (13)
Fadi	0.484 (7)	0.704 (4)	0.703 (4)	0.725 (6)	0.698 (5)	0.254 (20)	0.420 (5)
lys_acoruna	0.462 (9)	0.653 (9)	0.680 (7)	0.698 (9)	0.692 (6)	0.349 (10)	0.414 (8)
QiZhang	0.496 (5)	0.684 (5)	0.686 (5)	0.676 (10)	0.623 (11)	0.351 (8)	0.409 (10)
luxinyu	0.487 (6)	0.658 (8)	0.651 (9)	0.710 (7)	0.669 (8)	0.269 (19)	0.416 (7)
rafalposwiata	0.459 (10)	0.650 (10)	0.653 (8)	0.670 (11)	0.663 (9)	0.326 (13)	0.395 (12)
evanyfyang	0.213 (21)	0.635 (12)	0.639 (10)	0.703 (8)	0.642 (10)	0.350 (9)	0.449 (4)
robvanderg	0.366 (12)	0.648 (11)	0.605 (11)	0.632 (14)	0.614 (13)	0.296 (16)	0.344 (14)
psarangi	0.343 (15)	0.634 (13)	0.559 (12)	0.634 (13)	0.595 (14)	0.283 (17)	0.320 (17)
chx.dou	0.395 (11)	0.583 (14)	0.506 (13)	0.626 (15)	0.622 (12)	0.309 (14)	0.280 (20)

Table 3: Top 15 Results for crosslingual setting.

Crosslingual			
User	EN-ES	EN-CA	EN-EU
Cong666	0.644 (1)	0.643 (1)	0.632 (1)
colorful	0.618 (3)	0.562 (7)	0.584 (2)
gmorio	0.628 (2)	0.607 (3)	0.527 (4)
whu_stone	0.604 (5)	0.596 (4)	0.512 (7)
QiZhang	0.551 (10)	0.615 (2)	0.530 (3)
Fadi	0.589 (6)	0.593 (5)	0.516 (6)
hades_d	0.617 (4)	0.544 (10)	0.522 (5)
lys_acoruna	0.570 (7)	0.554 (8)	0.509 (8)
rafalposwiata	0.564 (8)	0.586 (6)	0.444 (12)
KE_AI	0.561 (9)	0.552 (9)	0.463 (11)
etms.kgp	0.542 (11)	0.506 (11)	0.431 (13)
jylong	0.375 (12)	0.474 (12)	0.504 (9)
ouzh	0.375 (12)	0.474 (12)	0.504 (9)
SPDB_Inn...	0.356 (13)	0.470 (13)	0.486 (10)
gerarld	0.321 (14)	0.269 (14)	0.303 (14)

et al., 2021) to use Graph F1 scores to evaluate structured sentiment extraction and classification.

In detail, Table 2 shows the results of the monolingual setting. For English, the top 15 participants lie between 62.6% and 76.0%, 30.9% and 44.7%, 28.0% and 49.4% F1 score on the OpeNer_{en}, MPQA, DS_{Unis} datasets. For the Spanish dataset, the top 15 F1 scores lie between 62.2% and 72.2%. The participants in the middle of the table are quite close to each other. For Catalan, F1 scores range from 58.3% to 72.8%. For the Basque language,

F1 scores range from 50.6% to 73.9% , with a gap of 20 points between first and last place. For the Basque dataset, F1 scores range from 50.6% to 73.9%, and most are in mid 60%. In this setting, we obtain ninth place based on the average score of all the datasets.

Table 3 shows the performance of crosslingual setting, which trains on English datasets and tests on the target languages including MultiBooked datasets and the OpeNER Spanish dataset. For Spanish, the top 15 F1 scores between 32.1% and 64.4%. For Catalan, F1 scores between 26.9% and 64.3%, which span a wide range. For the Basque dataset, the F1 scores range from 30.3% to 63.2%. Particularly, our model obtains the second and third places on the MultiB_{CA} and MultiB_{EU} languages.

5 Conclusions

In this paper, we propose a simple and unified model for both the monolingual and crosslingual structured sentiment analysis tasks. Different from the previous studies, we transform this task into an event extraction task. Mainly, we first design an expression extraction for extracting the expression, just like extracting the trigger words in the event extraction task. Then we predict the holder and target based on the extraction results of the previous step. The model performs well on seven datasets in five languages.

Acknowledgements

The authors wish to thank the reviewers for their helpful comments and suggestions. This research is funded by the Science and Technology Commission of Shanghai Municipality (No. 19511120200 and 21511100402). This research is also funded by the National Key Research and Development Program of China (No. 2021ZD0114002), the National Nature Science Foundation of China (No. 61906045), and Shanghai Science and Technology Innovation Action Plan International Cooperation project "Research on international multi language online learning platform and key technologies (No.20510780100)". The computation is performed in ECNU Multi-functional Platform for Innovation (001).

References

- Jeremy Barnes, Toni Badia, and Patrik Lambert. 2018. [MultiBooked: A corpus of Basque and Catalan hotel reviews annotated for aspect-level sentiment classification](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Jeremy Barnes, Robin Kurtz, Stephan Oepen, Lilja Øvrelid, and Erik Velldal. 2021. [Structured sentiment analysis as dependency graph parsing](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3387–3402, Online. Association for Computational Linguistics.
- Zhuang Chen and Tiejun Qian. 2020. Relation-aware collaborative learning for unified aspect-based sentiment analysis. In *Proceedings of ACL*, pages 3685–3694.
- Yejin Choi, Eric Breck, and Claire Cardie. 2006. Joint extraction of entities and relations for opinion recognition. In *Proceedings of the 2006 conference on empirical methods in natural language processing*, pages 431–439.
- Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. 2017. Modulating early visual processing by language. *Advances in Neural Information Processing Systems*, 30.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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 *Proceedings of ACL*, pages 504–515.
- Arzoo Katiyar and Claire Cardie. 2016. Investigating lstms for joint extraction of opinion entities and relations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 919–929.
- Robin Kurtz, Stephan Oepen, and Marco Kuhlmann. 2020. End-to-end negation resolution as graph parsing. In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 14–24.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.
- Lilja Øvrelid, Petter Mæhlum, Jeremy Barnes, and Erik Velldal. 2020. [A fine-grained sentiment dataset for Norwegian](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5025–5033, Marseille, France. European Language Resources Association.
- Eric S Tellez, Sabino Miranda-Jiménez, Mario Graff, Daniela Moctezuma, Ranyart R Suárez, and Oscar S Siordia. 2017. A simple approach to multilingual polarity classification in twitter. *Pattern Recognition Letters*, 94:68–74.
- Cigdem Toprak, Niklas Jakob, and Iryna Gurevych. 2010. [Sentence and expression level annotation of opinions in user-generated discourse](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 575–584, Uppsala, Sweden. Association for Computational Linguistics.
- Xiaojun Wan. 2008. Using bilingual knowledge and ensemble techniques for unsupervised chinese sentiment analysis. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 553–561.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation*, 39(2):165–210.
- Bishan Yang and Claire Cardie. 2012. Extracting opinion expressions with semi-markov conditional random fields. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1335–1345.
- Meishan Zhang, Qiansheng Wang, and Guohong Fu. 2019. End-to-end neural opinion extraction with a transition-based model. *Information Systems*, 80:56–63.

Jie Zhou, Jimmy Xiangji Huang, Qin Chen, Qinmin Vivian Hu, Tingting Wang, and Liang He. 2019. Deep learning for aspect-level sentiment classification: survey, vision, and challenges. *IEEE access*, 7:78454–78483.

Jie Zhou, Jimmy Xiangji Huang, Qinmin Vivian Hu, and Liang He. 2020a. Modeling multi-aspect relationship with joint learning for aspect-level sentiment classification. In *International Conference on Database Systems for Advanced Applications*, pages 786–802. Springer.

Jie Zhou, Jimmy Xiangji Huang, Qinmin Vivian Hu, and Liang He. 2020b. Sk-gcn: Modeling syntax and knowledge via graph convolutional network for aspect-level sentiment classification. *Knowledge-Based Systems*, 205:106292.