Learning Emotion-Aware Contextual Representations for Emotion-Cause Pair Extraction

Baopu Qiu[1,2] and Lin Shang[1,2(✉)]

[1] State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China
qiubaopu@smail.nju.edu.cn
[2] Department of Computer Science and Technology, Nanjing University, Nanjing 210023, China
shanglin@nju.edu.cn

Abstract. Emotion-Cause Pair Extraction (ECPE) focuses on analyzing emotions and its corresponding causes in a document. Two reasons have made ECPE a more challenging, but more applicable task than the previous Emotion-Cause Extraction (ECE) task: 1) an ECPE model needs to identify both emotions and their corresponding causes without the annotation of emotions. 2) the ECPE task involves finding causes for multiple emotions in a document, while ECE is for one emotion. However, existing ECPE methods fail to meet the second challenge, since they are evaluated on a dataset which exhibits a bias that nearly 90% of documents have only one emotion-cause pair. Thus, we reconstruct the dataset to better meet ECPE settings. We observe that previous SOTA approaches suffer from performance degradation in extracting multiple emotion-cause pairs due to the use of shared context encoder in the joint learning process. In this work, we propose a new pipelined approach that builds on two independent models with unshared context encoders, in which the emotion extraction model only provides input features for the cause extraction model. Experimental results demonstrate that our model can learn distinct contextual representations specific to each emotion, reaching state-of-the-art performance on both datasets and showing robustness in the analysis of more complex document context.

Keywords: Emotion-cause pair extraction · Emotion cause analysis · Representation learning

1 Introduction

Recently, the task about detecting the stimuli of emotions expressed in text has emerged in the area of text emotion analysis. Previous works focus on Emotion Cause Extraction (ECE), which has been proposed by [1] as a word-level

---

Our code and dataset are avaliable at https://github.com/Qbop981001/emotion-aware-eca.

sequence labeling problem. [2] re-formalized ECE as a clause-level classification problem of finding cause clauses for the given emotion. They released a Chinese dataset collected from SINA city news which has become the benchmark dataset for the ECE task followed by many works [3–7].

Some researchers [8] pointed out that ECE task suffers from two defects: 1) The emotion must be annotated in advance. 2) The goal of ECE neglects the fact that emotions and causes are mutually indicative. They developed the task to emotion-cause pair extraction (ECPE), in which emotion clauses and their corresponding cause clauses are extracted as pairs. To solve the problem, they proposed a two-step pipelined approach, while more recently, the task has been dominated by end-to-end systems that model emotion extraction and cause extraction jointly [9–11].

We re-investigate ECPE's motivation and observe another significant merit of ECPE over ECE, that is, the task involves analysis of more complex document context that contains multiple emotions, multiple causes and multiple semantic roles. As shown in Fig. 1, the example document is divided into *two* different samples in the ECE task since two emotion clauses are annotated. An ECE model takes each annotated emotion clause as input and find its corresponding causes, while an ECPE model takes the whole document as input and extract all possible emotion-cause pairs. Therefore, the model has to process richer, but more complex context information.

Unfortunately, the benchmark ECPE dataset, which previous works are evaluated on, fails to capture this merit. Only 10.23% of documents in this dataset have more than one emotion-cause pair, and only 6.63% of documents have more than one emotion clause. This is due to the fact that the dataset is originally designed for the previous ECE task and many different "documents" are actually excerpts of a same news article. Therefore, we reconstruct the dataset by merging documents with the same context to better meet ECPE settings.

We conduct experiments on the reconstructed dataset and find that current state-of-the-art ECPE methods suffers from severe performance degradation in extracting multiple emotion-cause pairs. We argue that this can be attributed to the shared contextual representations of input document, since these methods jointly perform the emotion extraction and cause extraction process using a shared context encoder. The entangled contextual representation may hinder the model from focusing on the proper part of context when finding causes for a specific emotion. For instance, the clauses c12 and c13 in Fig. 1, "failing to cure the disease while using that much money" is crucial in detecting the causal relationship between c14 and c15, but not relevant for c18 and c17.

To address this problem, we propose a new pipelined approach that builds on two independent context encoders trained separately, one for an emotion extraction model, and another for an emotion-oriented cause extraction model, with the fusion of emotion information merely at its input layer. Based on our experiments, we validate that the cause extraction model can learn emotion-aware contextual representations of the input document. Our approach reaches state-of-the-art in both datasets, and is significantly more effective in extracting

multiple emotion-cause pairs, meeting ECPE's motivation to analyze emotion causes in longer and more complex document context.

Our contributions can be summarized as follows:

– We realize another merit of the ECPE task over previous task: it enables the analysis of causes for multiple emotions, which usually occur in longer document context. We argue that current ECPE approaches fail to exploit this merit since the benchmark dataset is biased. We propose a strategy to reconstruct the dataset to better meet ECPE settings.
– We observe that existing models suffer from performance drop in extracting multiple emotion-cause pairs, and attribute this to the use of shared context encoder during the joint learning process. We propose a pipelined approach that learns two independent context encoders, with early fusion of emotion information at the input layer of the cause model to learn contextual representations specific to each extracted emotion.
– Experimental results on both datasets show that our approach achieves state-of-the-art performance in both datasets, and is more effective in extracting multiple emotion-cause pairs.
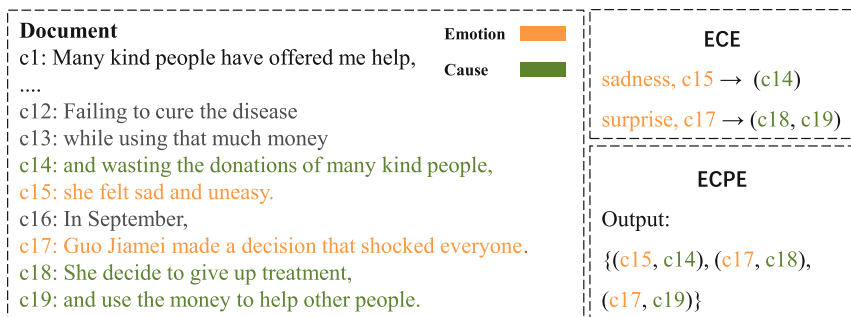
Fig. 1. An Example document from [2]'s dataset, translated from Chinese. Texts in orange and green denote the emotion clauses and cause clauses respectively. (Color figure online)

2   Related Work

Emotion Cause Extraction. [1] first proposed the emotion cause extraction task and released a small scale dataset. Early works adopted rule-based [14], machine-learning-based [15] methods to solve the task. Based on analysis of linguistic features in a Chinese dataset, researchers [16] have suggested that a clause may be the most proper unit for emotion cause analysis in Chinese. [2] reformalized the task as clause-level binary classification and released a benchmark corpus for the ECE task, followed by many works [3–6].

Emotion Cause Pair Extraction. [8] expanded the task to emotion-cause pair extraction and construct a benchmark ECPE corpus based the [2]'s dataset. [8] proposed a two-step pipeline approach to solve the task, of which the first step uses two component to extract all emotion clauses and cause clauses from the document, and attempts cartesian product to form all possible pairs. In the second step, the candidate pairs are fed into a filter to select emotion-cause pairs. More recently, most following works employ end-to-end models [10,17,20,21] with the belief that joint models capture interactions between subtasks and mitigate error propagation. Some of the models select the result from all possible pairs [9,12,19], and some others regard ECPE as a clause-level sequence labeling problem [13,18]. All these methods focus on the utilization of document context information and have employed various architectures and techniques such as graph convolutional network [19], graph attention network [11] and iterative synchronized multi-task learning [12].

Pipeline Approach vs. Joint Approach. The entity and relation joint extraction task, which involves extracting entities and their relations simultaneously, is a well-known task in information extraction, and is similar to the emotion-cause classifying and pairing process. Many existing works model entity extraction and relation classification jointly while some [22] argued that shared contextual representations during the joint learning process lead to sub-optimal results. They proposed a pipelined approach using separate encoders, and reached state-of-the-art performance.

3   Preliminary

3.1   Task Definition

Emotion Cause Extraction. Emotion Cause Extraction (ECE) has been defined as a clause-level classification task [2] aiming at extracting the corresponding cause of the annotated emotion in the context. Given a document $d = [c_1, ..., c_i, ..., c_{|d|}]$, where c_i is the ith clause in d, and an annotated emotion clause c^e, where $e \in E$,

$$E = \{happiness, sadness, disgust, fear, anger, surprise\} \qquad (1)$$

The goal of ECE is to find all the cause clauses of the given emotion clause as $\{c_1^{cau}, c_2^{cau}, ...\}$. Note that only one emotion occur in one sample, while there may be multiple causes corresponding to it.

Emotion-Cause Pair Extraction. [8] developed the ECE task to Emotion-Cause Pair Extraction (ECPE). Given a document $d = [c_1, ..., c_i, ..., c_{|d|}]$, the goal of ECPE is to extract a set of emotion-cause pairs

$$P = \{..., (c^{emo}, c^{cau}), ...\}$$

where c^{emo} is the emotion clause and c^{cau} is its corresponding cause clause. The ECPE task deals with finding multiple causes for multiple emotions in one document.

3.2   Dataset

Bias in the ECPE Benchmark Dataset. Based on SINA city news, [2] released an Chinese emotion cause corpus that has become the benchmark dataset for ECE research, which involves extraction of causes for one annotated emotion. Thus, a large proportion of documents in this dataset contain only one emotion clause, and documents with multiple emotions are split to different samples. Following researchers [8] also use this dataset as the benchmark for ECPE. They merged samples with same text content into one document since in the ECPE task, every document corresponds to one sample.

As Table 1 shows, we can observe that the dataset exhibits a bias that only 10.23% of documents contain multiple emotion-cause pairs, and only 6.63% of documents have more than one emotion. The bias prevents researchers from knowing the performance of their models on multiple emotion-cause pair extraction, which can be regarded as a significant merit of ECPE over previous tasks. We also discover that many different "documents" are actually excerpts from the *same* original news report. During the construction of original ECE dataset, people may only select clauses surrounding the annotated emotion and ignore long-range clauses. This reduces the task's difficulty at the expense of applicability, since in real word, documents such as news article and literary work usually contain multiple emotions, belonging to multiple semantic roles.

Table 1. Statistics of the original dataset and reconstructed dataset

Item	Original dataset	Reconstructed dataset
Doc. total number	1945	1679
Doc. with one emotion cause pair	1746	1348
Doc. with two emotion cause pairs	177	246
Doc. with more than two pairs	22	85
Doc. with more than one emotion	129	295

Dataset Reconstruction Strategy
We manually find all such documents and merge them into one document to rebuild the ECPE benchmark dataset. As Table 1 shows, 17.57% of documents in the reconstructed dataset have multiple emotions while 19.71% of documents have multiple emotion-cause pairs. Our merging strategy reduces the total number of documents and produces longer documents with increased complexity. Figure 2 displays the comparision of the number of clauses in a document between the original and reconstructed dataset. As is shown, the documents in the reconstructed dataset have more clauses and thus more complex document structure. In fact, 37.42% of emotion-cause pairs are located in documents with multiple pairs, indicating that documents in the reconstructed dataset are closer to real-world scenarios.
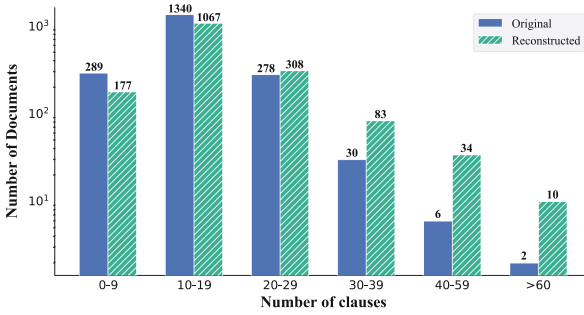
Fig. 2. Statistics of the document length in both datasets.

4 Methodology

As Fig. 3 shows, our approach consists of two independent models, an emotion extraction model and a cause extraction model. We build both of our models on BERT [23] as context encoders, with an multi-label output layer. The emotion extraction model first takes the whole document as input and extract all possible emotion clauses. Then the extracted emotion clauses will be used *one by one* for fusing emotion information at the input layer of the cause extraction model, which we refer to as an emotion-oriented cause extraction model. We will explain the details of both models below and clarify the usage of emotion information as well as document context information in our approach.
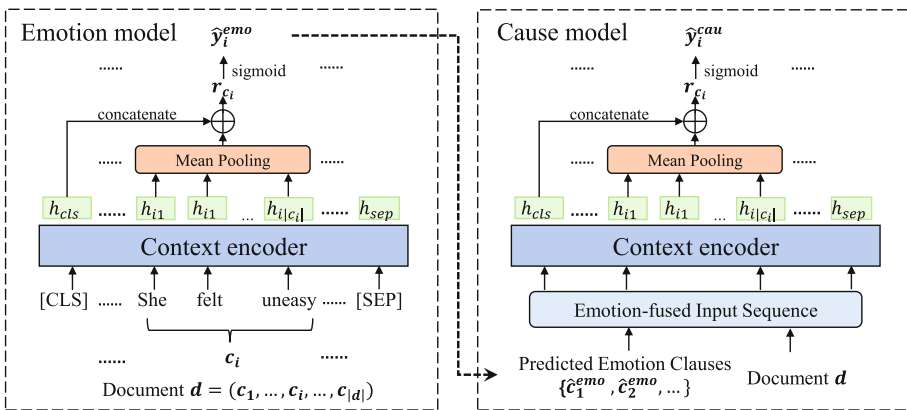
4.1 Emotion Extraction Model

Fig. 3. Model architecture.

Given a document $d = [c_1, ..., c_i, ..., c_{|d|}]$, the model takes it as the input of a pre-trained encoder to obtain a sequence of hidden states denoted by

$$\mathbf{H_D} = (h_{[CLS]}, \mathbf{x_{c_1}}, ..., \mathbf{x_{c_i}}, ..., \mathbf{x_{c_{|d|}}}, h_{[SEP]}) \tag{2}$$

where $\mathbf{x_{c_i}} = (h_{i1}, ..., h_{ij}, ...h_{i|c_i|})$, $h_{ij} \in R^H$ is the output hidden state of j^{th} token in i^{th} clause and $|c_i|$ denotes the number of tokens in i^{th} clause. Then we apply mean pooling to build the representation of each clause, which is defined as:

$$\mathbf{h_{c_i}} = \frac{1}{|c_i|} \sum_{j=1}^{|c_i|} h_{ij} \tag{3}$$

Finally we concatenate the clause representation $\mathbf{h_{c_i}}$ with $[CLS]$ token's output hidden state, $h_{[CLS]}$, as the input of an output layer to predict the probability of the clause being an emotion clause

$$\mathbf{r_{c_i}} = [\mathbf{h_{c_i}}, h_{[CLS]}] \tag{4}$$

$$\hat{y}_i^{emo} = \sigma(w_{emo}^T \mathbf{r_{c_i}} + b_{emo}) \tag{5}$$

where $w_{emo} \in R^{H \times 1}$ and b_{emo} are parameters of the output layer with sigmoid function $\sigma(\cdot)$.
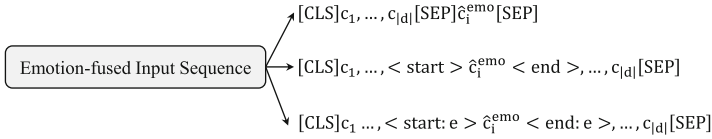
Context Information. In order to examine the impact of context information in the emotion extraction process, we also implement a standard BERT-based sentence classification model in which each single clause is taken as the input without leveraging the context. The details are explained in Sect. 5.3.

4.2 Emotion-Oriented Cause Extraction Model

In the two-step model proposed by [8], there is an cause extraction component that extracts potential cause clauses in a document at first. We found the performance of this component unsatisfying since it ignores the fact that the identification of certain cause clause depends on its corresponding emotion clause. In our approach, we do *not* perform cause extraction *solely*, and instead conduct emotion-oriented cause extraction.

As Fig. 3 shows, the architecture of cause model is very similar to our emotion extraction model, and the only difference lies in the input: we fuse emotion information into the input sequence through three strategies.

Fusing Emotion Information. Previous works have attached importance to the use of emotion information in cause extraction [11,12,21]. However, all of these models use a shared LSTM layer or pre-trained encoder for contextual representations in emotion extraction and cause extraction. We argue that shared context encoders fail to capture proper contextual information for a specific emotion clause, leading to sub-optimal results in extracting multiple emotion-cause pairs in one document. Therefore, we propose strategies to fuse specific emotion information at the input layer of the cause model, as displayed in Fig. 4.

$$[\text{CLS}]c_1, \ldots, c_{|d|}[\text{SEP}]\hat{c}_i^{\text{emo}}[\text{SEP}]$$

Emotion-fused Input Sequence $\longrightarrow$ $[\text{CLS}]c_1, \ldots, <\text{start}> \hat{c}_i^{\text{emo}} <\text{end}>, \ldots, c_{|d|}[\text{SEP}]$

$$[\text{CLS}]c_1 \ldots, <\text{start}:e> \hat{c}_i^{\text{emo}} <\text{end}:e>, \ldots, c_{|d|}[\text{SEP}]$$

Fig. 4. Emotion-fusing strategies. From top to the bottom, the strategy is denoted by EmotionPrompt, UntypedMarker and TypedMarker in the following expeiments. $e \in E$ is the type of emotion defined in Eq. 1.

The most direct method is to concatenate each predicted emotion clause and its document context as the input. We also attempt to integrate emotion information through extra marker tokens at the start and end position of the predicted emotion clause, and further consider adding the emotion type explicitly. However, since we cannot obtain satisfying results for emotion type classification, we only compare the upper bound of its effectiveness with other emotion-fusing strategies using the ground truth emotion type label. The details of comparative results are elaborated in Sect. 5.3.

4.3 Training and Inference

For both models, we fine-tune the pre-trained encoder using task-specific training objectives. Given a document d, we compute the loss for both models by:

$$L = -\frac{1}{|d|} \sum_{i=1}^{|d|} H(\hat{y}_i, y) \tag{6}$$

where $|d|$ is the number of clauses in the document, $H(\cdot)$ is the binary cross-entropy loss function, $\hat{y}_i$ is $\hat{y}_i^{emo}$ defined by Eq. 5 in the emotion model and $\hat{y}_i^{cau}$ in the cause model, while y is the ground truth label of the clause.

During training, the two models are trained separately, and we use ground truth emotion labels to fuse emotion information in the cause extraction model. During inference, we first use the emotion model to extract emotion clauses in each document and fed each predicted emotion clause into the cause extraction model to generate emotion-aware contextual representations.

5 Experiments

5.1 Evaluation

For evaluation metrics, precison, recall and F1 defined in [8] are used. Most of previous ECPE approaches also evaluate their models on two subtasks: emotion extraction and cause extraction, yet we do this only for emotion extraction since our approach do not perform cause extraction solely.

Table 2. Comparative results of existing models and our approach. For fair comparison, if a model has an implementation based on BERT, we report the BERT-based results, and use †to mark the models that are not BERT-based

Model	Original dataset			Reconstructed dataset			Multiple pairs		
	P(%)	R(%)	F1(%)	P(%)	R(%)	F1(%)	P(%)	R(%)	F1(%)
Indep†[8]	68.32	50.82	58.18	–	–	–	–	–	–
Inter-CE†	69.02	51.35	59.01	–	–	–	–	–	–
Inter-EC†	67.21	57.05	61.28	–	–	–	–	–	–
USL [18]†	71.49	62.79	66.86	–	–	–	–	–	–
SLSN [17]†	68.36	62.91	65.45	–	–	–	–	–	–
LAE-MANN [21]	71.10	60.70	65.50	–	–	–	–	–	–
Tagging [13]	72.43	63.66	67.76	–	–	–	–	–	–
ECPE-2D [9]	72.92	65.44	68.89	–	–	–	–	–	–
RANK-CP [11]	71.19	76.30	73.60	77.89	49.90	60.69	76.88	53.96	63.15
PairGCN [19]	76.92	67.91	72.02	–	–	–	–	–	–
ECPE-MLL [12]	77.00	72.35	74.52	68.46	67.06	67.65	68.84	55.90	61.55
Ours-EmotionPrompt	**77.83**	76.01	**76.81**	71.05	**74.83**	72.85	69.62	61.85	65.19
Ours-UntypedMarker	76.27	75.83	75.96	73.78	73.30	**73.50**	70.94	**62.80**	**66.50**
w/o emotion-fusing	69.70	71.10	70.36	61.95	64.74	59.73	41.06	42.27	41.14

5.2 Experimental Settings

We implement our approach based on `Pytorch` and `Transformers` and use `bert-base-chinese` as the base encoder. For both models, we set the random seed to 42 and use Adam optimizer for training. The learning rate is 2e−5, warmup ratio is 0.1, and threshold of the multi-label output layer is 0.5. In the experiments, we follow previous works [8,9,11,13,17–19,21] to perform 10-fold cross validation and use the same data split of the original dataset.

5.3 Results and Analysis

Comparative Approaches. Most of existing ECPE works are joint models using shared context encoders, except from **Indep**, **Inter-CE** and **Inter-EC**, the three variants of the two-step pipelined models proposed by [8] that serve as the baseline. **Rank-CP** [11] and **ECPE-MLL** [12] are the two previous state-of-the-art approaches and thus we evaluate and compare the performance of these two approaches with ours on the reconstructed dataset as well as on a subset of documents that *only* contain more than one emotion-cause pair, denoted by "Multiple pairs." It should be noted that the pair selection process of the **Rank-CP** model relies on a sentiment lexicon, which may be inflexible in a wider range of application scenarios.

Main Results. Table 2 displays the comparative results. As is shown, our approach achieves state-of-the-art performance in both datasets. Our approach with `EmotionPrompt` and `UnTypedMarker` achieves an absolute F1 improvement of
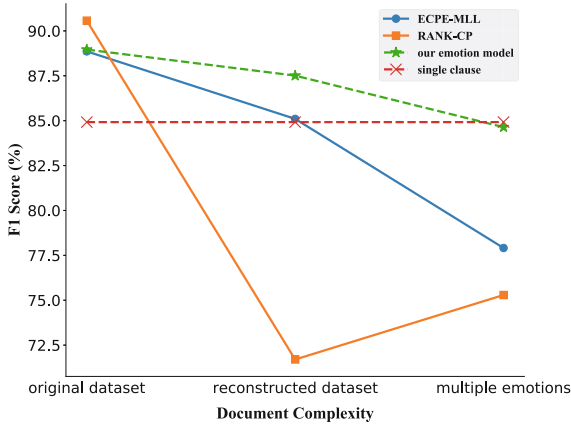
Fig. 5. Results on emotion extraction.

2.29% and 1.44% respectively over the best previous work [12] on the original dataset, and an improvement of 5.20% and 5.85% respectively over [12] on the reconstructed dataset. For the comparison of pipelined approaches, our approach outperform the baseline Inter-EC by 15.53% and 14.68% respectively.

Results on Extracting Multiple Emotion-Cause Pairs. The results show that our approach with UntypedMarker outperforms [11]'s previous work by an absolute F1 of 3.35% on extracting multiple pairs. The performance of [11]'s model increases on multiple pairs mainly because they apply the sentiment lexicon to filter candidate emotion-pairs and tend to select fewer pairs, resulting in high precision rate and low recall rate.

Importance of Emotion-Fusing. In the previous part, we attach importance to contextual representations specific to each emotion clause and fusion of emotion information at the input layer of the cause extraction model. Above results show that both emotion-fusing strategies achieve convincing results, and in order to further validate the impact of emotion-fusing, we conduct ablation experiments by removing emotion information in the cause model.

As shown in Table 2, we can observe a clear gap between our models and the model without fusion of emotion features, especially in the reconstructed dataset and on multiple emotion-cause pairs extraction. Since the classification of an emotion-cause heavily depends on the emotion it corresponds to, it is almost meaningless to perform cause extraction without emotion information, with the decline of 18.59% F1 score in extracting multiple pairs.

Results on Emotion Extraction One motivation of joint approaches in ECPE is that the performance of emotion extraction can also be improved by cause information provided during the joint training process. Indeed, we can observe from Fig. 5 that joint models outperform our model on the original dataset. In

Table 3. Comparative results of the upper bound of emotion-aware cause Extraction

Strategy	P(%)	R(%)	F1(%)
Inter-EC Bound [8]	76.10	70.84	73.28
LAE-ECE [21]	80.80	79.90	80.30
UntypedMarker	84.68	83.54	84.09
TypedMarker	85.44	**84.43**	84.92
EmotionPrompt	**85.99**	83.98	**84.95**

the reconstructed dataset, however, their performance exhibits a clear decline, and even worse in extracting multiple emotions from one document. We assume that cause information obtained via joint training does bring some benefits, but as document complexity grows, shared encoders in joint models fail to capture proper information from entangled context, and the entangled contextual representations provide more noise than benefits for the model. Comparative results between our model and single classification model demonstrate that context information encoded in $h_{[CLS]}$ benefits less as document complexity increases.

Upper Bound of Emotion-Aware Cause Extraction. We consider using emotion type information explicitly. We test the upper bound of emotion-aware cause extraction using ground truth emotion type label for each emotion clause and compare the results between different emotion-fusing strategies and other ECPE methods which also report their upper bound results.

As shown in Table 3, we observe the benefits brought by emotion type between UntypedMarker and TypedMarker, while EmotionPrompt obtains the best F1 score, indicating that it may be better to integrate emotion information through emotional text. Futhermore, there are couple of documents that exceed the max input length of BERT. We split such documents to different parts in the experiments, but text markers cannot be used if the emotion clause is located in another part of an document. Thus, for future works, we suggest the use of EmotionPrompt, which is more flexible, as the emotion-fusing strategy.

5.4   Case Study: Capture Emotion-aware Document Context

In this subsection, we discuss a specific document example in the dataset, which contains 25 clauses, 3 emotion-pairs and 2 different emotion (`c7`,`c10`). Part of the document is listed below:

..., c4: the six members of the family live on the two or three thousand Yuan Gong earns from working every month, c5: with so many children, c6: and poor conditions at home, c7: Mr.Gong was very sad. c8: What bothered him more was that, c9: because of over childbirth, c10: his child's registered permanent residence could not be solved.
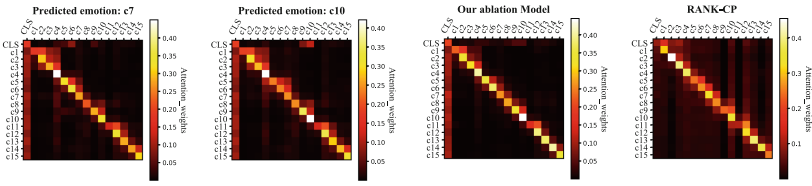
(a) Our EmotionPrompt model fused (b) RANK-CP [11] and our ablation
with different emotions c7,c10 at the model without emotion fusing.
input layer.

Fig. 6. Document attention heatmap produced by ECPE models for the example in
Sect. 5.4. In order to make the figure more intuitive, we only select attention of the
first 15 clauses.

As Fig. 6 shows, the fusion of emotion enables our model EmotionPrompt
to capture proper document context information. c4 is important in finding
causes for both emotions, since it clarifies the background of the document.
When finding causes for c7, c5 and c6, which elaborate the concrete condition
of Mr.Gong of being poor, are important. For emotion clause c10, c9 and c10
explain why Mr.Gong is "bothered."

As Fig. 6 shows, both of our ablation model and Rank-CP model fail to
capture emotion-aware context, either due to the lack of emotion-guided input,
or the entangled representation obtained via joint training.

6   Conclusion

In this paper, we realize another significant merit of the ECPE task, which
is extracting multiple emotion-cause pairs from longer context, and find that
the existing ECPE works fail to capture this merit due to bias in the bench-
mark dataset they are evaluated on. We reconstruct the dataset and conduct
experiments on both datasets, observing that previous SOTA works of ECPE
suffer from a performance drop on multiple emotion-cause pair extraction due
to the use of shared context encoders. To address the problem, we present a
simple but effective approach that builds on two independent context encoders.
Experimental results demonstrate that our approach can learn contextual rep-
resentations specific to each emotion and reaches state-of-the-art performance
on both datasets, while showing robustness in extracting multiple emotion-cause
pairs among more complex document context.

References

1. Lee, S.Y., Chen, Y., Huang, C.R.: A text-driven rule-based system for emotion cause detection. In: Proceedings of the NAACL HLT Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, pp. 45–53 (2010)
2. Gui, L., Wu, D., Xu, R., Lu, Q., Zhou, Y.: Event-driven emotion cause extraction with corpus construction. In: EMNLP, pp. 1639–1649 (2016)
3. Gui, L., Hu, J., He, Y., Xu, R., Lu, Q., Du, J.: A question answering approach for emotion cause extraction. In: EMNLP, pp. 1593–1602 (2017)
4. Li, X., Song, K., Feng, S., Wang, D., Zhang, Y.: A co-attention neural network model for emotion cause analysis with emotional context awareness. In: EMNLP, pp. 4752–57 (2018)
5. Xia, R., Zhang, M., Ding, Z.: RTHN: A RNN-transformer hierarchical network for emotion cause extraction. arXiv preprint arXiv:1906.01236 (2019)
6. Fan, C., et al.: A knowledge regularized hierarchical approach for emotion cause analysis. In: EMNLP, pp. 5614–5624 (2019)
7. Yan, H., Gui, L., Pergola, G., He, Y.: Position bias mitigation: a knowledge-aware graph model for emotion cause extraction. In: ACL, pp. 3364–3375 (2021)
8. Xia, R,. Ding, Z.: Emotion-cause pair extraction: a new task to emotion analysis in texts. In: ACL, pp. 1003–1012 (2019)
9. Ding, Z., Xia, R., Yu, J.: ECPE-2D: emotion-cause pair extraction based on joint two-dimensional representation, interaction and prediction. In: ACL, pp. 3161–70(2020)
10. Fan, C., Yuan, C., Du, J., Gui, L., Yang, M., Xu, R.: Transition-based directed graph construction for emotion-cause pair extraction. In: ACL, pp. 3707–3717 (2020)
11. Wei, P., Zhao, J., Mao, W.: Effective inter-clause modeling for end-to-end emotion-cause pair extraction. In: ACL, pp. 3171–3181 (2020)
12. Ding, Z., Xia, R., Yu, J.: End-to-end emotion-cause pair extraction based on sliding window multi-label learning. In: EMNLP, pp. 3574–3583 (2020)
13. Yuan, C., Fan, C., Bao, J., Xu, R.: Emotion-cause pair extraction as sequence labeling based on a novel tagging scheme. In: EMNLP, pp. 3568–3573 (2020)
14. Gao, K., Xu, H., Wang, J.: A rule-based approach to emotion cause detection for Chinese micro-blogs. Expert Syst. Appl. **42**(9), 4517–28 (2015)
15. Ghazi, D., Inkpen, D., Szpakowicz, S.: Detecting emotion stimuli in emotion-bearing sentences. In: Gelbukh, A. (ed.) CICLing 2015. LNCS, vol. 9042, pp. 152–165. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-18117-2_12
16. Chen, Y., Lee, S.Y., Li, S., Huang, C.R.: Emotion cause detection with linguistic constructions. In: COLING, pp. 179–187 (2010)
17. Cheng, Z., Jiang, Z., Yin, Y., Yu, H., Gu, Q.: A symmetric local search network for emotion-cause pair extraction. In: COLING, pp. 139–149 (2020)
18. Chen, X., Li, Q., Wang, J.: A unified sequence labeling model for emotion cause pair extraction. In: COLING, pp. 208–218 (2020)
19. Chen, Y., Hou, W., Li, S., Wu, C., Zhang, X.: End-to-end emotion-cause pair extraction with graph convolutional network. In: COLING, pp. 198–207 (2020)
20. Fan, R., Wang, Y., He, T.: An end-to-end multi-task learning network with scope controller for emotion-cause pair extraction. In: Zhu, X., Zhang, M., Hong, Yu., He, R. (eds.) NLPCC 2020. LNCS (LNAI), vol. 12430, pp. 764–776. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-60450-9_60

21. Tang, H., Ji, D., Zhou, Q.: Joint multi-level attentional model for emotion detection and emotion-cause pair extraction. Neurocomputing **409**, 329–340 (2020)
22. Zhong, Z., Chen, D.: A frustratingly easy approach for entity and relation extraction. In: NAACL, pp. 50–61 (2021)
23. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: NAACL, pp. 4171–4186 (2019)