

An End-to-End Network for Emotion-Cause Pair Extraction

Aaditya Singh^{*†} Shreeshail Hingane[†] Saim Wani[†] Ashutosh Modi

Indian Institute of Technology Kanpur (IIT Kanpur)

{aadisin, shreesh, saim}@iitk.ac.in

ashutoshm@cse.iitk.ac.in

Abstract

The task of Emotion-Cause Pair Extraction (ECPE) aims to extract all potential clause-pairs of emotions and their corresponding causes in a document. Unlike the more well-studied task of Emotion Cause Extraction (ECE), ECPE does not require the emotion clauses to be provided as annotations. Previous works on ECPE have either followed a multi-stage approach where emotion extraction, cause extraction, and pairing are done independently or use complex architectures to resolve its limitations. In this paper, we propose an end-to-end model for the ECPE task. Due to the unavailability of an English language ECPE corpus, we adapt the NTCIR-13 ECE corpus and establish a baseline for the ECPE task on this dataset. On this dataset, the proposed method produces significant performance improvements ($\sim 6.5\%$ increase in F1 score) over the multi-stage approach and achieves comparable performance to the state of the art methods.

1 Introduction

There have been several works on emotions prediction from the text (Alswaidan and Menai, 2020; Witon et al., 2018) as well as generating emotion oriented texts (Ghosh et al., 2017; Colombo et al., 2019; Modi et al., 2020; Goswamy et al., 2020). However, recently the focus has also shifted to finding out the underlying cause(s) that lead to the emotion expressed in the text. In this respect, Gui et al. (2016) proposed the Emotion Cause Extraction (ECE), a task aimed at detecting the cause behind a given emotion annotation. The task is defined as a clause level classification problem. The text is divided at the clause level and the task is to detect the clause containing the cause, given the clause containing the emotion. However, the

applicability of models solving the ECE problem is limited by the fact that emotion annotations are required at test time. More recently, Xia and Ding (2019) introduced the Emotion-Cause Pair Extraction (ECPE) task i.e. extracting all possible emotion-cause clause pairs in a document with no emotion annotations. Thus, ECPE opens up avenues for applications of real-time sentiment-cause analysis in tweets and product reviews. ECPE builds on the existing and well studied ECE task. Figure 1 shows an example with ground truth annotations.

Xia and Ding (2019) use a two-stage architecture to extract potential emotion-cause clauses. In Stage 1, the model extracts a set of emotion clauses and a set of cause clauses (not mutually exclusive) from the document. In Stage 2, it performs emotion-cause pairing and filtering, i.e. eliminating pairs that the model predicts as an invalid emotion-cause pair. However, this fails to fully capture the mutual dependence between emotion and cause clauses since clause extraction happens in isolation from the pairing step. Thus, the model is never optimized using the overall task as the objective. Also, certain emotion clauses are likely not to be detected without the corresponding cause clauses as the context for that emotion. Recent methods such as Ding et al. (2020a) and Ding et al. (2020b) use complex encoder and classifier architectures to resolve these limitations of the multi-stage method.

In this paper, we propose an end-to-end model to explicitly demonstrate the effectiveness of joint training on the ECPE task. The proposed model attempts to take into account the mutual interdependence between emotion and cause clauses. Based on the benchmark English-language corpus used in the ECE task of the NTCIR-13 workshop (Gao et al., 2017), we evaluate our approach on this dataset after adapting it for the ECPE task. We demonstrate that the proposed approach works sig-

* Currently affiliated with IBM Research

† Authors equally contributed

Clause 1:	Adele arrived at her apartment late in the afternoon after a long day of work.
Clause 2:	She was still furious with her husband for not remembering her 40th birthday.
Clause 3:	As soon as she unlocked the door, she gasped with surprise;
Clause 4:	Mikhael and Harriet had organized a huge party for her.

Figure 1: An example document. The above example contains two emotion-cause pairs. Clause 2 is an emotion clause (*furious*) and is also the corresponding cause clause (*for not remembering her 40th birthday*). Clause 3 is an emotion clause (*surprise*) and Clause 4 is its corresponding cause clause (*organized a huge party for her*).

nificantly better than the multi-stage method and achieves comparable performance to the state of the art methods. We also show that when used for the ECE task by providing ground truth emotion annotations, our model beats the state of the art performance of ECE models on the introduced corpus. We provide the dataset and implementations of our models via GitHub¹.

2 Related Work

The problem of Emotion Cause Extraction (ECE) has been studied extensively over the past decade. ECE was initially proposed as a word-level sequence detection problem in Lee et al. (2010). Attempts to solve this task focused on either classical machine learning techniques (Ghazi et al., 2015), or on rule-based methods (Neviarouskaya and Aono, 2013; Gao et al., 2015). Subsequently, the problem was reframed as a clause-level classification problem (Chen et al., 2010) and the Chinese-language dataset introduced by Gui et al. (2016) has since become the benchmark dataset for ECE and the task has been an active area of research (Xia et al., 2019; Gui et al., 2017; Yu et al., 2019; Li et al., 2018, 2019; Fan et al., 2019).

However, the main limitation of ECE remains that it requires emotion annotations even during test time, which severely limits the applicability of ECE models. To address this, Xia and Ding (2019) introduced a new task called emotion-cause pair extraction (ECPE), which extracts both emotion and its cause without requiring the emotion annotation. They demonstrated the results of their two-stage architecture on the benchmark Chinese language ECE corpus (Gui et al., 2016). Following their work, several works have been proposed to address the limitations of the two-stage architecture (Ding et al. (2020a), Ding et al. (2020b), Fan et al. (2020), Yuan et al. (2020), Cheng et al. (2020), Chen et al. (2020)). In order to explore the corpus further and to encourage future work from a broader commu-

nity, we use an English language ECE corpus. (see section 5).

3 Task Definition

Formally, a document consists of text that is segmented into an ordered set of clauses $D = [c_1, c_2, \dots, c_d]$ and the ECPE task aims to extract a set of emotion-cause pairs $P = \{\dots, (c_i, c_j), \dots\}$ ($c_i, c_j \in D$), where c_i is an emotion clause and c_j is the corresponding cause clause. In the ECE task, we are additionally given the annotations of emotion clauses and the goal is to detect (one or more) clauses containing the cause for each emotion clause.

4 Approach

We propose an end-to-end emotion cause pairs extraction model (Figure 2), henceforth referred to as E2E-PE_{EXT_E} (refer to section 6 for the naming convention). The model takes an entire document as its input and computes, for each ordered pair of clauses (c_i, c_j) , the probability of being a potential emotion-cause pair. To facilitate the learning of suitable clause representations required for this *primary* task, we train the model on two other *auxiliary* tasks: Emotion Detection and Cause Detection.

We propose a hierarchical architecture. Word level representations are used to obtain clausal representations (via BiLSTM) and clause level representations are further *contextualized* using another BiLSTM network. The resulting contextualized clause representations are then used for the classification task. Let w_i^j represent the vector representation of the j^{th} word in the i^{th} clause. Each clause c_i in the document d is passed through a *word-level* encoder (BiLSTM + Attention (Bahdanau et al., 2015)) to obtain the clause representation s_i . The clause embeddings are then fed into two separate *clause-level* encoders (Emotion-Encoder and Cause-Encoder) each of which corresponds, respectively, to the two auxiliary tasks. The purpose of the clause-level encoders is to help learn

¹<https://github.com/Aaditya-Singh/E2E-ECPE>




Figure 2: End-to-End network (E2E-PExtE)

contextualized clause representations by incorporating context from the neighboring clauses in the document. For each clause c_i , we obtain its contextualized representations \mathbf{r}_i^e and \mathbf{r}_i^c by passing it through a BiLSTM network.

For the clause c_i , its contextual representations \mathbf{r}_i^e and \mathbf{r}_i^c are then used to predict whether the clause is an emotion-clause and a cause-clause respectively, i.e.,

$$\hat{\mathbf{y}}_i^e = \text{softmax}(\mathbf{W}^e * \mathbf{r}_i^e + \mathbf{b}^e); \quad (1)$$

$$\hat{\mathbf{y}}_i^c = \text{softmax}(\mathbf{W}^c * \mathbf{r}_i^c + \mathbf{b}^c) \quad (2)$$

As observed by Xia and Ding (2019), we also noticed that performance on the two auxiliary tasks could be improved if done in an interactive manner rather than independently. Hence, the Cause-Encoder also makes use of the corresponding emotion-detection prediction $\hat{\mathbf{y}}_i^e$, when generating \mathbf{r}_i^c (Figure 2).

For the primary task, every ordered pair (c_i, c_j) is represented by concatenating \mathbf{r}_i^e , \mathbf{r}_j^c and \mathbf{pe}_{ij} , wherein \mathbf{pe}_{ij} is the positional embedding vector representing the relative positioning between the two clauses i, j in the document (Shaw et al., 2018).

The primary task is solved by passing this pair-representation through a fully-connected neural network to get the pair-predictions.

$$\mathbf{r}_{ij}^p = [\mathbf{r}_i^e \oplus \mathbf{r}_j^c \oplus \mathbf{pe}_{ij}]; \quad (3)$$

$$\mathbf{h}_{ij}^p = \text{ReLU}(\mathbf{W}^{p1} * \mathbf{r}_{ij}^p + \mathbf{b}^{p1}); \quad (4)$$

$$\hat{\mathbf{y}}_{ij}^p = \text{softmax}(\mathbf{W}^{p2} * \mathbf{h}_{ij}^p + \mathbf{b}^{p2}) \quad (5)$$

To train the end-to-end model, loss function is set as the weighted sum of loss on the primary task as well as the two auxiliary tasks i.e.

$L^{\text{total}} = \lambda_c * L^c + \lambda_e * L^e + \lambda_p * L^p$, where L^e , L^c , and L^p are cross-entropy errors for emotion, cause and pair predictions respectively. Further, $L^p = L^{\text{pos}} + \text{loss_weight} * L^{\text{neg}}$, where L^{pos} and L^{neg} are the errors attributed to positive and negative examples respectively. We use hyperparameter loss_weight to scale down L^{neg} , since there are far more negative examples than positive ones in the primary pairing task.

5 Corpus/Data Description

We adapt an existing Emotion-Cause Extraction (ECE) (Fan et al., 2019; Li et al., 2019) corpus for evaluating our proposed models (as well as the architectures proposed in previous work). The corpus was introduced in the NTCIR-13 Workshop (Gao et al., 2017) for the ECE challenge.

The corpus consists of 2843 documents taken from several English novels. Each document is annotated with the following information: i) emotion-cause pairs present in the document, that is, the set of emotion clauses and their corresponding cause clauses; ii) emotion category of each clause; and iii) the keyword within the clause denoting the labeled emotion. We do not use the emotion category or the keyword during training of either ECE or ECPE tasks only the emotion-cause pairs are used. At test time, none of the annotations are used for the ECPE task. For ECE task, emotion annotation is provided at test time and the model predicts the corresponding cause clauses. 80%-10%-10% splits are used for training, validation and testing. 10 such randomly generated splits are used to get statistically significant results, and the average results are reported.

	Emotion Extraction			Cause Extraction			Pair Extraction		
	Precision	Recall	F1 Score	Precision	Recall	F1 Score	Precision	Recall	F1 Score
ECPE 2-stage (Xia and Ding, 2019)	67.41	71.60	69.40	60.39	47.34	53.01	46.94	41.02	43.67
ECPE-2D (BERT) (Ding et al., 2020a)	74.35	69.68	71.89	64.91	53.53	58.55	60.49	43.84	50.73
ECPE-MLL (ISML-6) (Ding et al., 2020b)	75.46	69.96	72.55	63.50	59.19	61.10	59.26	45.30	51.21
E2E-PExt _E (Ours)	71.63	67.49	69.43	66.36	43.75	52.26	51.34	49.29	50.17
RHNN (Fan et al., 2019)	-	-	-	-	-	-	69.01	52.67	59.75
E2E-CExt (Ours)	-	-	-	-	-	-	65.21	66.18	65.63

Table 1: Experimental results comparing our models with the already existing benchmarks. The top half compares our model for ECPE (E2E-PExt_E) against the existing benchmarks. The bottom half compares existing ECE benchmark on this dataset against our model (E2E-CExt). Note that the Pair Extraction task in ECPE with true-emotion provided reduces to the Cause Extraction task of ECE. The results on RHNN and E2E-CExt are only for the primary task, since in the ECE setting, there are no auxiliary tasks. The evaluation metrics are same as the ones used in previous works.

6 Experiments and Results

Naming Scheme for Model Variants:

E2E: End to End.

PExt/CExt/EExt: Pair Extraction/ Cause Extraction/ Emotion Extraction.

Subscript E/C: represents how the auxiliary tasks are solved interactively (emotion predictions used for cause detection or vice versa).

For the purpose of evaluation on our dataset, we reproduced the two-stage model: ECPE 2-stage (Xia and Ding, 2019). We also adapted two models that achieve state of the art performance on the ECPE task: ECPE-2D (BERT) (Ding et al., 2020a) and ECPE-MLL (ISML-6) (Ding et al., 2020b) and compared them against our model: E2E-PExt_E. The model is trained for 15 epochs using Adam optimizer (Kingma and Ba, 2014). The learning rate and batch size were set to 0.005 and 32 respectively. Model weights and biases were initialized by sampling from a uniform distribution $U(-0.10, 0.10)$. GloVe word embeddings (Pennington et al., 2014) of 200 dimension are used. For regularization, we set the dropout rate to 0.8 for word embeddings and L2 weight decay of $1e-5$ over softmax parameters. We set $\lambda_c : \lambda_e : \lambda_p = 1 : 1 : 2.5$. The values chosen through grid search on the hyperparameter space reflect the higher importance of the primary pair detection task compared to the auxiliary tasks. To obtain better positional embeddings which encode the relative positioning between clauses, we trained randomly initialized embeddings after setting the clipping distance (Shaw et al., 2018) to 10 all clauses that have a distance of 10 or more between them have the same positional embedding.

We use the same evaluation metrics (precision, recall and F1-score), as used in the past work on ECE and ECPE tasks. Following Xia and Ding

(2019), the metric definitions are defined as:

$$P = \frac{\#correct_pairs}{\#proposed_pairs}$$

$$R = \frac{\#correct_pairs}{\#annotated_pairs}$$

$$F1 = \frac{2 * P * R}{P + R}$$

where,

$\#proposed_pairs$ = no. of emotion-cause pairs predicted by model

$\#correct_pairs$ = number of emotion-cause pairs predicted *correctly* by the model

$\#annotated_pairs$ = total number of actual emotion-cause pairs in the data

$P, R, F1$ for the two auxiliary classification tasks (emotion-detection and cause-detection) have the usual definition (see Gui et al. (2016)).

The results are summarized in Table 1. Our model E2E-PExt_E outperforms ECPE 2-stage on the task of emotion-cause pair extraction by a significant margin of 6.5%. This explicitly demonstrates that an end-to-end model works much better since it leverages the mutual dependence between the emotion and cause clauses. As shown in Table 1, our model achieves comparable performance to the highly parameterized and complex models ECPE-2D (BERT) and ECPE-MLL (ISML-6) (which either leverage a pre-trained BERT (Devlin et al., 2019) and 2D Transformer (Vaswani et al., 2017) or iterative BiLSTM encoder). To further demonstrate this point, we compare the number of trainable parameters across models in Table 2.

Method	Trainable parameters
E2E-PE _{Ext_E}	790,257
ECPE-2D (BERT)	1,064,886
ECPE-MLL (ISML-6)	6,370,452

Table 2: Comparison of trainable parameters of our model (E2E-PE_{Ext_E}) with the state-of-the-art models (ECPE-2D (BERT) and ECPE-MLL (ISML-6)). We achieve comparable performance with these models with fewer parameters and simpler architecture.

We also evaluate our model on the ECE task. For this, we use a variant of E2E-PE_{Ext_E} i.e. E2E-CE_{Ext}, so that it utilizes the emotion annotations. Specifically, we incorporate the knowledge of emotion annotations by incorporating them into the Cause-Encoder as well as the Pair-Prediction-Module and show that this improves performance in both the primary pair prediction task as well as the auxiliary cause detection task (appendix section A). E2E-CE_{Ext} outperforms the state of the art model RHNN: (Fan et al., 2019) on the ECE task. This is indicative of the generalization capability of our model and further demonstrates that performance on ECPE can be enhanced with improvements in the quality of emotion predictions.

6.1 Ablation Experiments

We analyzed the effect of different components of the model on performance via ablation experiments. We present some notable results below. More results are presented in the Appendix.

Positional Embeddings: For finding the extent to which positional embeddings affect the performance, we train E2E-PE_{Ext_E} without positional embeddings. This resulted in a slight drop in validation F1 score from 51.34 to 50.74, which suggests that the network is robust to withstanding the loss of distance information at the clause-level.

Loss Weighting: To handle the problem of data imbalance, we varied loss weight. We observed (Figure 3) that assigning less weight to the negative examples leads to more predicted positives, and hence a better recall but worse precision.

7 Conclusion

In this paper, we demonstrated that a simple end-to-end model can achieve competitive performance on the ECPE task by leveraging the inherent correlation between emotions and their causes and optimizing directly on the overall objective. We




Figure 3: Precision, Recall and F1 Score as a function of weight assigned to negative examples.

also showed that a variant of our model which further uses emotion annotations, outperforms the previously best performing model on the ECE task, thereby showing its applicability to variety of related tasks involving emotion analysis. In future, we plan on developing a larger benchmark English-language dataset for the ECPE task. We plan to explore other model-architectures which we expect will help us learn richer representations for causality detection in clause-pairs.

References

- Nourah Alswaidan and Mohamed El Bachir Menai. 2020. A survey of state-of-the-art approaches for emotion recognition in text. *Knowledge and Information Systems*, pages 1–51.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Xinhong Chen, Qing Li, and Jianping Wang. 2020. A unified sequence labeling model for emotion cause pair extraction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 208–218.
- Ying Chen, Sophia Yat Mei Lee, Shoushan Li, and Churen Huang. 2010. Emotion cause detection with linguistic constructions. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 179–187.
- Zifeng Cheng, Zhiwei Jiang, Yafeng Yin, Hua Yu, and Qing Gu. 2020. A symmetric local search network for emotion-cause pair extraction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 139–149.
- Pierre Colombo, Wojciech Witon, Ashutosh Modi, James Kennedy, and Mubbasir Kapadia. 2019.

- Affect-driven dialog generation.** 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 3734–3743, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Zixiang Ding, Rui Xia, and Jianfei Yu. 2020a. Ecpe-2d: Emotion-cause pair extraction based on joint two-dimensional representation, interaction and prediction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3161–3170.
- Zixiang Ding, Rui Xia, and Jianfei Yu. 2020b. End-to-end emotion-cause pair extraction based on sliding window multi-label learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3574–3583.
- Chuang Fan, Hongyu Yan, Jiachen Du, Lin Gui, Lidong Bing, Min Yang, Ruifeng Xu, and Ruibin Mao. 2019. **A knowledge regularized hierarchical approach for emotion cause analysis.** 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)*, pages 5614–5624, Hong Kong, China. Association for Computational Linguistics.
- Chuang Fan, Chaofa Yuan, Jiachen Du, Lin Gui, Min Yang, and Ruifeng Xu. 2020. Transition-based directed graph construction for emotion-cause pair extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3707–3717.
- Kai Gao, Hua Xu, and Jiushuo Wang. 2015. **A rule-based approach to emotion cause detection for chinese micro-blogs.** *Expert Syst. Appl.*, 42(9):45174528.
- Qinghong Gao, Gui Lin, Yulan He, Jiannan Hu, Qin Lu, Ruifeng Xu, and Kam-Fai Wong. 2017. Overview of ntcir-13 eca task. In *Proceedings of the NTCIR-13 Conference*.
- Diman Ghazi, Diana Inkpen, and Stan Szpakowicz. 2015. Detecting emotion stimuli in emotion-bearing sentences. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 152–165. Springer.
- Sayan Ghosh, Mathieu Chollet, Eugene Laksana, Louis-Philippe Morency, and Stefan Scherer. 2017. Affect-lm: A neural language model for customizable affective text generation. *arXiv preprint arXiv:1704.06851*.
- Tushar Goswamy, Ishika Singh, Ahsan Barkati, and Ashutosh Modi. 2020. **Adapting a language model for controlled affective text generation.** In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2787–2801, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Lin Gui, Jiannan Hu, Yulan He, Ruifeng Xu, Qin Lu, and Jiachen Du. 2017. **A question answering approach for emotion cause extraction.** In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1593–1602, Copenhagen, Denmark. Association for Computational Linguistics.
- Lin Gui, Dongyin Wu, Ruifeng Xu, Qin Lu, and Yu Zhou. 2016. **Event-driven emotion cause extraction with corpus construction.** In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1639–1649, Austin, Texas. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Sophia Yat Mei Lee, Ying Chen, and Chu-Ren Huang. 2010. **A text-driven rule-based system for emotion cause detection.** In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 45–53, Los Angeles, CA. Association for Computational Linguistics.
- Xiangju Li, Shi Feng, Daling Wang, and Yifei Zhang. 2019. Context-aware emotion cause analysis with multi-attention-based neural network. *Knowl. Based Syst.*, 174:205–218.
- Xiangju Li, Kaisong Song, Shi Feng, Daling Wang, and Yifei Zhang. 2018. **A co-attention neural network model for emotion cause analysis with emotional context awareness.** In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4752–4757, Brussels, Belgium. Association for Computational Linguistics.
- Ashutosh Modi, Mubbasir Kapadia, Douglas A. Fidaio, James R. Kennedy, Wojciech Witon, and Pierre Colombo. 2020. Affect-driven dialog generation. *US Patent 10,818,312*.
- Alena Neviarouskaya and Masaki Aono. 2013. Extracting causes of emotions from text. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 932–936.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.

Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.

Wojciech Witon, Pierre Colombo, Ashutosh Modi, and Mubbasir Kapadia. 2018. [Disney at IEST 2018: Predicting emotions using an ensemble](#). In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 248–253, Brussels, Belgium. Association for Computational Linguistics.

Rui Xia and Zixiang Ding. 2019. [Emotion-cause pair extraction: A new task to emotion analysis in texts](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1003–1012, Florence, Italy. Association for Computational Linguistics.

Rui Xia, Mengran Zhang, and Zixiang Ding. 2019. Rthn: a rnn-transformer hierarchical network for emotion cause extraction. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 5285–5291. AAAI Press.

X. Yu, W. Rong, Z. Zhang, Y. Ouyang, and Z. Xiong. 2019. Multiple level hierarchical network-based clause selection for emotion cause extraction. *IEEE Access*, 7:9071–9079.

Chaofa Yuan, Chuang Fan, Jianzhu Bao, and Ruifeng Xu. 2020. Emotion-cause pair extraction as sequence labeling based on a novel tagging scheme. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3568–3573.

Appendix

A Model Variants

We present a variant of $E2E-PExt_E$ called $E2E-PExt_C$ (Figure 4). Here, we feed the clause embeddings s_i into a clause-level BiLSTM to obtain context-aware clause encodings r_i^c which is fed into a softmax layer to obtain cause predictions y_i^c . We concatenate the cause predictions with the clause embeddings, $s_i \oplus y_i^c$, and feed them into another clause-level BiLSTM to obtain emotion representations r_i^e , which are fed into another softmax layer to obtain the emotion predictions y_i^e . The pair prediction network remains the same as described for $E2E-PExt_E$. For finding the extent

to which emotion labels can help in improving pair predictions, we present a variant of $E2E-PExt_E$ called $E2E-CExt$. We use the true emotion labels instead of emotion predictions y_i^e to obtain the context-aware clause encodings r_i^c . We also concatenate them with the input of pair-prediction network $r_i^e \oplus r_j^c \oplus pe_{ij}$ to make full use of the additional knowledge of emotion labels. Similarly, the corresponding variant of $E2E-PExt_C$ which utilizes true cause labels is called $E2E-EEExt$. The results of model variation are shown in Table 3. Here, the pair prediction network consists of a single fully connected layer. After comparing the performance of $E2E-PExt_C$ and $E2E-EEExt$ we conclude that huge improvements in performance can be achieved if the quality of cause predictions is improved.

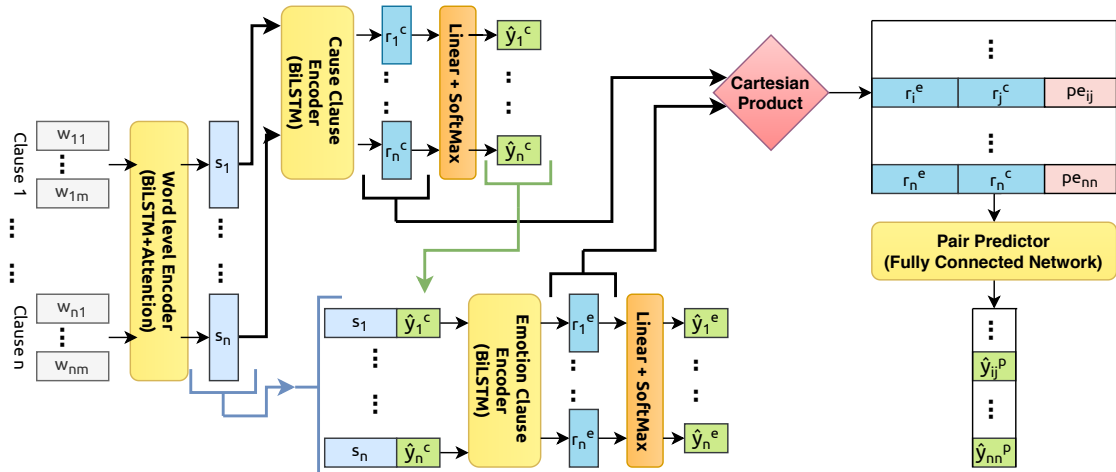


Figure 4: End-to-End network variant (E2E-PExt_c)

	Emotion Extraction			Cause Extraction			Pair Extraction		
	Precision	Recall	F1 Score	Precision	Recall	F1 Score	Precision	Recall	F1 Score
E2E-PExt _c	71.70	66.77	69.10	63.75	42.50	50.42	48.88	48.22	48.37
E2E-EEExt	81.74	71.95	76.47	100.00	100.00	100.00	74.17	78.63	76.30

Table 3: Results of Model variants using precision, recall, and F1-score on the ECPE task and the two sub-tasks.