

# LyS at SemEval-2024 Task 3: An Early Prototype for End-to-End Multimodal Emotion Linking as Graph-Based Parsing

Ana Ezquerro and David Vilares

Universidade da Coruña, CITIC

Departamento de Ciencias de la Computación y Tecnologías de la Información

Campus de Elviña s/n, 15071

A Coruña, Spain

{ana.ezquerro, david.vilares}@udc.es

## Abstract

This paper describes our participation in SemEval 2024 Task 3, which focused on Multimodal Emotion Cause Analysis in Conversations. We developed an early prototype for an end-to-end system that uses graph-based methods from dependency parsing to identify causal emotion relations in multi-party conversations. Our model comprises a neural transformer-based encoder for contextualizing multimodal conversation data and a graph-based decoder for generating the adjacency matrix scores of the causal graph. We ranked 7th out of 15 valid and official submissions for Subtask 1, using textual inputs only. We also discuss our participation in Subtask 2 during post-evaluation using multi-modal inputs.

## 1 Introduction

SemEval 2024 Task 3 focused on Multimodal Emotion Cause Analysis in Conversations (Wang et al., 2024). Figure 1 shows an example provided by the organizers to illustrate the task. Two subtasks were proposed: Subtask 1, which uses only textual inputs, and Subtask 2, which allows for the consideration of video and audio processing as well.

The shared task is timely given the recent success of multimodal architectures combining computer vision (Redmon et al., 2016; Wang et al., 2023b), natural language processing (Devlin et al., 2019; Beltagy et al., 2020), and speech processing (Gong et al., 2021; Radford et al., 2022) advancements. In the particular context of multimodal emotion analysis, the task builds on top of previous work such as recognizing the triggered emotions as a classification task (Alhuzali and Ananiadou, 2021; Zheng et al., 2023) or predicting complex cause-effect relations between speakers (Wei et al., 2020; Ding et al., 2020). For the particular case of the shared task, the dataset - centered in English - relies on (Wang et al., 2023a) and provides text, image and audio inputs.




Figure 1: Example taken from the official website of the SemEval Task 3 - [https://nustm.github.io/SemEval-2024\\_ECAC/](https://nustm.github.io/SemEval-2024_ECAC/). The goal of the task consists of predicting (i) the emotion associated to each utterance within the conversation, (ii) the cause-effect relations that trigger the emotions between utterances and (iii) the associated span in the cause utterance.

We had time and resources only to build a textual model for official participation in Subtask 1. We validated some multimodal baseline approaches using vision and audio inputs, but the computational resources required to fine-tune text and video data were beyond our range, so we participated in Subtask 2 only during post-evaluation. In what follows, we describe our approach. The implementation of our early prototype can be found at <https://github.com/anaezquerro/semEval24-task3>.

**Contribution** We propose an end-to-end multimodal prototype based on a large multimodal encoder to contextualize text, image and audio inputs with a graph-based decoder to model the cause-effect relations between triggered emotions within multi-party conversations. The large encoder joins pretrained architectures in text (Devlin et al., 2019), image (Dosovitskiy et al., 2021) and audio (Baevski et al., 2020) modalities, while the decoder is adapted from the graph-based approaches in semantic parsing (Dozat and Manning, 2018). The model is trained end-to-end.

## 2 Background

**Multimodal Emotion Cause Analysis** A number of datasets collecting multi-party conversations

(Poria et al., 2019; Chatterjee et al., 2019; Firdaus et al., 2020) have been published to train and test multimodal neural architectures. Simpler configurations involve recognizing the speaker emotion at each utterance - this task is commonly known as Emotion Recognition (ER) (Poria et al., 2019) - while others require a deeper level of understanding to model interactions and causal relations between speakers - Emotion-Cause Pair Extraction (ECPE) (Xia and Ding, 2019). The most common approaches follow an encoder-decoder neural architecture where the encoder is conformed by multiple modules - one module per input modality (text, image and/or audio) - and produces an inner representation at utterance level; and the decoder accepts the encoder outputs as inputs and returns a suitable output adapted to the specifications of the targeted task. In the context of Multimodal ER, Nguyen et al. (2023) proposed a GCN-based decoder to capture temporal relations (Schlichtkrull et al., 2017), while Dutta and Ganapathy (2024) used cross-attention to fusion the input modalities and a final classification layer to predict the targeted emotions. Approaches in ECPE require an extra effort to represent and model causal information: Wei et al. (2020) scored all possible utterance tuples to predict the most probable list of emotion-cause pairs. Other authors, like Chen et al. (2020) and Fan et al. (2020), represented the emotion-cause pairs as a labeled graph between utterances and tried to predict the set of causal edges using a GCN or a transition-based system, respectively. The SemEval 2024 Task 3 joins the recognition and causal extraction tasks and challenges a system able to both model speaker emotions and elicit relations.

**Graph-based decoding** For structured prediction tasks, such as dependency parsing, graph-based approaches are a standard for computing output syntactic representations (McDonald, 2006; Martins et al., 2013). Particularly, Dozat and Manning (2017) introduced a classifier that computes a head and dependent representation for each token and then uses two biaffine classifiers: one computes a score for each pair of tokens to determine the most likely head, and the other determines the label for each head-dependent token pair. We will also build upon a biaffine graph-based parser: we will frame the task as predicting a dependency graph, where utterances are the nodes and emotions are dependency labels between pairs of utterances.

### 3 System Overview

Our system consists of two modules: a large pre-trained encoder and a graph-based decoder (see Figure 2). It can add extra input channels into the encoder without requiring any adjustments to the decoder, so the same decoder is used for both tasks, while the encoder is adapted to incorporate text-only (Subtask 1) or multimodal (Subtask 2) inputs.




Figure 2: High-level architecture of our system. The encoder takes as input the sequence of  $m$  utterances of a given conversation and returns a unique vector representation for each utterance. The decoder uses the utterance embedding matrix to apply the affine attention product in the decoder, obtain the scores of the adjacent matrix and return the predicted sequence of emotions and the cause relations between utterances.

Let  $C = (U_1, \dots, U_m)$  be a conversation of  $m$  utterances, where each utterance  $U_i = \{W_i, s_i, \varepsilon_i\}$  is defined by (i) a sequence of words  $W_i = (w_1^{(i)}, \dots, w_{\ell|w,i|}^{(i)})^1$ , (ii) an active speaker  $s_i$  and (iii) a triggered emotion  $\varepsilon_i \in \mathcal{E}^2$ . The set of cause-pair relations between utterances can be represented as a directed labeled graph  $G = (\mathcal{U}, \mathcal{R})$  where  $\mathcal{U} = (U_1, \dots, U_m)$  is the sequence of utterances of the conversation assuming the role of the nodes of the graph and  $\mathcal{R} = \{U_i \xrightarrow{\varepsilon_j} U_j, i, j \in [1, m]\}$  is the set of emotion-cause relations between an arbitrary cause utterance  $U_i$  and its corresponding effect  $U_j$ . Thus, the task can be cast as the estimation of the adjacent matrix of  $G$ , similarly to syntactic (Ji et al., 2019) and semantic dependency (Dozat and Manning, 2018) parsing. Adapting algorithms from parsing to model emotion-cause relations between utterances has also been explored by other authors, such as Fan et al. (2020), who instead explored a transition-based strategy.

<sup>1</sup>From now on, we denote as  $\ell|\cdot, i|$  the length of the  $i$ -th in a sequence  $\cdot$ , so  $\ell|w, i|$  denotes the length of the  $W_i$ . Table 2 summarizes the notation used in this paper.

<sup>2</sup>The set of emotions are described in Wang et al. (2023a).

### 3.1 Textual Extraction

The first subtask draws from only textual information to predict the adjacent matrix of  $G$  with a span that covers the specific words from  $U_i$  that trigger the emotion  $\varepsilon_j$  in the cause relation  $U_i \xrightarrow{\varepsilon_j} U_j$ .

**Textual encoder** Figure 3 illustrates our encoder. Given the sequence of utterances  $(U_1, \dots, U_m)$ , we encoded with BERT the batched sequence of utterances where each word sequence was preceded by the CLS token (Devlin et al., 2019). For each  $U_i$ , we select the CLS embedding ( $\mathbf{u}_i$ ) from the contextualized embedding matrix  $\mathbf{W}_i = (\mathbf{u}_i, \mathbf{w}_1, \dots, \mathbf{w}_{\ell|w,i|})$ , which is assumed to have information of the whole sentence. The CLS embedding matrix  $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_m)$  was passed as input to the decoder module and the word embeddings were reserved for the span attention module.




Figure 3: High level representation of the textual encoder. The input (1) is the matrix of stacked token vectors of each utterance. The last hidden states of BERT are used as word embeddings (2) and the special CLS tokens are used as utterance embeddings (3). The effect embeddings (4) - a partial representation from the decoder - are taken as input to the span module with the contextualized BERT embeddings.

**Graph-based decoder** Figure 4 shows the forward-pass of the graph-based decoder from the encoder output of Figure 3. To produce an adjacent matrix  $\mathbf{G}$  of dimensions  $m \times m$ , where each position  $(i, j)$  represents the probability of a causal relation from  $U_i$  (cause) to  $U_j$  (effect), the first biaffine module uses a trainable matrix  $\mathcal{W}_G \in \mathbb{R}^{d_G \times d_G}$  and maps  $\mathbf{U}$  using two feed-forward networks to a cause ( $\mathbf{C}$ ) and an effect ( $\mathbf{E}$ ) representation. By projecting the original BERT embeddings to two different representations,  $\mathbf{u}_i \sim (\mathbf{c}_i, \mathbf{e}_i)$ , the decoder learns different contributions for the same utterance depending on the role. The affine product is




Figure 4: Graph-based decoder. The utterance embeddings (1) are projected to different representations (2, 3, 5, 6) using four feed-forward networks to flexibly represent utterance embeddings. The scores of the adjacent matrix and the probability tensor are computed with the affine attention product.

defined as  $\mathbf{G} = \mathbf{E} \cdot \mathcal{W} \cdot \mathbf{C}^\top$ . The second biaffine module uses a trainable tensor  $\mathcal{W}_\varepsilon \in \mathbb{R}^{d_G \times |\mathcal{E}| \times d_G}$  to predict the probabilities of triggered emotions between cause-effect utterances.

**Span Attention module** To maintain the end-to-end prediction while learning the span associated to each relation  $U_i \rightarrow U_j$ , we created a binary tensor  $\mathbf{S} = (\mathbf{S}_1 \cdots \mathbf{S}_m)$  of dimensions  $m \times m \times \max_{i=1, \dots, m} \{\ell|w, i|\}$ <sup>3</sup> to specify if a word  $w_k \in W_i$  of  $U_i$  is included in the span that triggers an emotion in  $U_j$ . To compute each  $\mathbf{S}_i$ , the matrix of word embeddings ( $\mathbf{W}_i$ ) of the utterance  $U_i$  is passed through a One-Head Attention module (see Figure 5), where  $\mathbf{W}_i$  acts as the query matrix and  $\mathbf{E}$  as the key and value matrices, so  $\mathbf{S}_i = \Phi(\text{softmax}(\mathbf{W}_i \cdot \mathbf{E}^\top) \cdot \mathbf{E})$ , where  $\Phi$  is a feed-forward network to project the embedding dimension to a unique binary value.

**Encoding speaker information** The dataset includes information about the active speakers in each utterance. A first approach to use this information as input would be concatenating the speaker embeddings to the sequence of utterances. However, this might lead to some issues: the model could assume that there is some inner dependency between triggered emotions and the characters in the conversation. This might be true in some cases, but it can also lead to biases, and there is still the challenge of modeling infrequent and unknown characters. To deal with this, we encoded a conversation  $C$  with speakers  $s_1, \dots, s_m$  using relative

<sup>3</sup>Note that each matrix  $\mathbf{S}_i$  has dimensions  $m \times \max_{i=1, \dots, m} \{\ell|w, i|\}$  and is associated to a *cause* utterance.




Figure 5: Span Attention module adapted from Vaswani et al. (2017). The tensor of word embeddings ( $\mathbf{W}_1 \cdots \mathbf{W}_m$ ) from the encoder (Figure 3) and the effect contextualizations ( $\mathbf{E}$ ) from the decoder (Figure 4) are passed to the attention product using each  $\mathbf{W}_i$  as *key* and  $\mathbf{E}$  as *value* matrices.

positional embeddings. For instance, the sequence (Chandler, Phoebe, Monica, Chandler, Phoebe) in Figure 1 would be encoded as (0, 1, 2, 0, 1).

### 3.2 Multimodal Extraction

The second subtask adds a short video representation to each utterance, so  $U_i$  in a conversation  $C = (U_1, \dots, U_m)$  is now a tuple of five different elements  $U_i = \{W_i, s_i, \varepsilon_i, \mathbf{X}_i, \mathbf{a}_i\}$ . The last two added items encode the image and audio: (i)  $\mathbf{X}_i = (\mathbf{x}_1^{(i)}, \dots, \mathbf{x}_{\ell[x,i]}^{(i)})$  is the sequence of frames of the input video, where each frame is an image<sup>4</sup> tensor of dimensions  $h \times w \times 3$  and (ii)  $\mathbf{a}_i$  is the sampled audio signal of arbitrary length.

**Image encoding** We relied on a Transformer-based architecture (Ma et al., 2022; Zheng et al., 2023) to contextualize input images. While recent studies have proposed adaptations of the Vision Transformer and 3-dimensional convolutions that capture temporal correlations between sequences of frames for video classification (Arnab et al., 2021; Ni et al., 2022), our experiments were constrained by our resource limitations, preventing us from using these pretrained architectures. Hence, for our multimodal baseline we opted for the the smallest version of the Vision Transformer (ViT) model (Dosovitskiy et al., 2021) pretrained on the Facial Emotion Recognition dataset (Goodfellow et al., 2013)<sup>5</sup> to contextualize a small fraction of sampled frames<sup>6</sup>, and incorporated an LSTM-based module to derive a unique image representation for each utterance. From an image batch  $\mathbf{X}_i$ , each image

<sup>4</sup>All frames are RGB images, being the majority resolution  $720 \times 1280$ .

<sup>5</sup><https://huggingface.co/trpakov/vit-face-expression>.

<sup>6</sup>For our experiments we used 5 interleaved frames per video, although a lower sampling rate can be considered depending on the computational capabilities.

$\mathbf{x}_k^{(i)} \in \mathbb{R}^{h \times w \times 3}$  was passed to the ViT base model to recover the output of the last hidden layer and introduce it as input to the LSTM module to recover a final representation for  $U_i$ .

**Audio encoding** For our multimodal system we used the hidden contextualizations of the base version of wav2vec 2.0 (Baevski et al., 2020)<sup>7</sup>. Given a raw audio ( $\mathbf{a}_i$ ) of an utterance  $U_i$ , the encoder of wav2vec 2.0 returns a sequence of hidden states that we summarized with an additional trainable LSTM layer to retrieve a unique vector that contains the audio information.

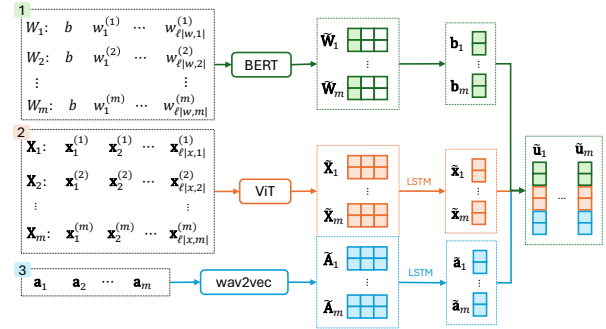


Figure 6: Multimodal encoder for Subtask 2.

**Model fine-tuning** The multimodal encoder (§3.2) uses three pretrained architectures to contextualize individual utterances and passes to the decoder the concatenation of the three unimodal representations (Figure 6). We chose to fine-tune only BERT during training together with the rest of the network. This was based on our empirical observation of superior results when learning from text compared to image and audio data. We entrusted the learning of audiovisual data to the LSTM learnable module within the encoder, presuming an accurate initial contextualization from wav2vec 2.0 and ViT pretrained on FER-2013.

### 3.3 Post-processing

Our end-to-end system directly recovers the predicted emotion-cause relations in a single post-processing step that linearly operates with the output tensors of the decoder. For the first subtask, the decoder returns (i) the adjacent matrix  $\mathbf{G} \in \mathbb{R}^{m \times m}$ , (ii) the labeled adjacent matrix  $\overline{\mathbf{G}} \in \mathbb{R}^{m \times m \times |\mathcal{E}|}$  and (iii) the span scores  $\mathbf{S} \in \mathbb{R}^{m \times m \times \ell_{\max}|w,i|}$ . As Dozat and Manning (2017), each arc  $U_i \rightarrow U_j$  is predicted by thresholding  $\mathbf{G}$ , and, once the arcs are predicted, the tensor  $\overline{\mathbf{G}}$  determines the label

<sup>7</sup><https://huggingface.co/facebook/wav2vec2-base-960h>.

(emotion) associated to each arc. Since our formalization (§3.1) associates a given utterance to an unique emotion, we leveraged the scores of  $\overline{\mathbf{G}}$  by the cause utterances and return the emotion with highest score. Finally, to produce a continuous span for each score vector  $\mathbf{s}_{ij}$ , we considered the leftmost and rightmost elements of  $\mathbf{s}_{ij}$  higher than a fixed threshold.

ST-1	$\mathbf{P}_s^w$	$\mathbf{R}_s^w$	$\dagger \mathbf{F}_s^w$	$\mathbf{P}_p^w$	$\mathbf{R}_p^w$	$\mathbf{F}_p^w$
BERT <sub>400</sub>	10.19	5.46	7.01	21.64	15.09	17.33
BERT <sub>600</sub>	12.61	7.43	9.32	22.06	15.2	17.95
BERT <sub>800</sub>	14.89	7.36	9.75	22.13	23.25	15.32

ST-2	$\mathbf{P}^w$	$\mathbf{R}^w$	$\dagger \mathbf{F}^w$
BERT	27.49,	17.62,	20.43
+ViT	22.38	22.72	22.17
+w2v	28.4	20.01	23.36
+w2v+ViT	23.37	7.62	11.49

Table 1: Evaluation of our prototype with different multimodal configurations. Precision (P), recall (R) and F-Score (F) measured the weighted average across the eight emotions of the dataset (superscript  $w$  denotes that the measure is weighted) and for the first subtask the span performance is considered with strict correctness (subscript  $s$ ) or overlapping (subscript  $p$ ). The symbol  $\dagger$  remarks the reference metric for each subtask.

## 4 Experiments

**Validation** The annotated dataset contains 1 375 multi-party conversations with a total of 13 619 utterances (Wang et al., 2023a). Although an unbiased estimation of the performance of our system would require validating the trained architecture using all available annotated data, our time and resources limitations prevented us from conducting k-fold cross-validation. Instead, we partitioned a 15% of the annotated dataset as our development set. The specific split used will be available with the accompanying code to replicate our findings.

**Evaluation** We use the official metrics<sup>8</sup>: the weighted strict F-Score for the Subtask 1 and the weighted F-Score for the Subtask 2.

**Hyperparameter configuration** Our computational limitations prevented us from exhaustively searching the optimal hyperparameters for our system. We conducted some tests varying the pre-

trained text encoder<sup>9</sup>, model dimension, gradient descent algorithm and learning rate and adding or removing the speaker module. We maintained in all experiments some regularization techniques (such as dropout in the hidden layers and gradient norm clipping) to avoid over-fitting. Our final configuration uses AdamW optimizer (Loshchilov and Hutter, 2019) with learning rate of  $10^{-6}$  and is trained during one hundred epochs with early stopping on the validation set.

## 5 Results

Table 1 presents the performance of our system for both subtasks. For the first subtask, we investigated various embedding sizes of the Biaffine decoder while concurrently fine-tuning the largest version of BERT<sup>10</sup>. For the second subtask, we conducted experiments using different types of inputs to evaluate their impact. These included: (i) using only text-based inputs, (ii) adding audio data, (iii) incorporating visual data through frames, and (iv) leveraging all available multimodal inputs together. For approaches (i), (ii) and (iii), only BERT was fine-tuned, whereas for approach (iv), all pretrained weights were frozen. These weights solely served to contextualize input information, with the learning process confined to the decoder component.

Our top-performing model for the first subtask achieved a validation score of 9.75 and ranked in the evaluation set in 7th position among 15 participants with 6.77 points. We observed a slight performance improvement by increasing the hidden dimension of the decoder. Thus, considering the expansion of decoder layers could improve the performance. It is worth noting the significant impact of span prediction on the model performance: the proportional results consistently outperform strict metrics. Removing span prediction while retaining only text inputs results in a notable increase in F-Score (20.43 points for the second subtask), indicating the crucial role of span prediction in model learning. Furthermore, we noticed that there was a consistent delay in the alignment between recall and precision metrics, with precision consistently exceeding recall by more than 5 points across all approaches. This suggests that our system tends to adopt a conservative behavior, avoiding the number

<sup>8</sup>[https://github.com/NUSTM/SemEval-2024\\_ECAC/tree/main/CodaLab/evaluation](https://github.com/NUSTM/SemEval-2024_ECAC/tree/main/CodaLab/evaluation).

<sup>9</sup>We performed some experiments using all the versions of BERT, (Devlin et al., 2019), RoBERTa (Liu et al., 2019) and XLM-RoBERTa (Conneau et al., 2020) and selected the best-performing textual encoder (BERT-large).

<sup>10</sup><https://huggingface.co/google-bert/bert-large-cased>

of false cause emotion predictions.

The best validation performance for the second subtask is achieved through the integration of text and audio, yielding a score of 23.36 points in the weighted F-Score. Using image data also improves the text-only baseline, though unexpectedly lags behind the audio model. It is important to note that these two approaches are not directly comparable due to differences in their data inputs: the text and image model only considers a fixed number of sampled frames, suggesting that providing more image data (ideally, the full sequence of frames) could potentially yield a better performance that surpasses the audio-based approach. Unfortunately, we could not fine-tune BERT with the full multimodal encoder, so we were restricted to projecting the multimodal inputs to their respective contextualizations, and relying on the trainable weights of the decoder to optimize the full architecture. The results prove the importance of, at least, fine-tuning the text encoder: the F-Score only reaches 11.25 points, whereas the text finetuned baseline nearly doubles its performance with 20.43 points, highlighting the insufficient context of the original pretrained BERT embeddings to address this task.

Once the post-evaluation period concluded, we upload an experimental submission of our best multimodal system to the official competition. We obtained 15.32 points in the weighted F-Score, positioning our baseline in the 13th place out of 18 participants.

## 6 Conclusion

We proposed a graph-based prototype for the analysis emotion-cause analysis in conversations. Given the limited preparation time, we only submitted official results for Subtask 1 (text-only), but also report post-evaluation results for Subtask 2 (multimodal). The task required predicting several aspects of the conversation: (i) the emotion associated with each utterance, (ii) the cause-effect relationships triggering these emotions between utterances, and (iii) the specific span within the cause utterance responsible for the emotion. We achieved 7th place out of 15 valid submissions for Subtask 1, a promising outcome considering the time and resource constraints we had to prepare the task. Yet, our results make us optimistic about exploring future research avenues to enhance our system and study lighter approaches that can perform competitively. As future work, we aim to experiment with

smaller and distilled models to encode textual, visual, and audio inputs, enabling us to fine-tune the full model cheaply.

## Acknowledgments

This work has received supported by Grant GAP (PID2022-139308OA-I00) funded by MCIN/AEI/10.13039/501100011033/ and by ERDF, EU; the European Research Council (ERC), under the Horizon Europe research and innovation programme (SALSA, grant agreement No 101100615); Grant SCANNER-UDC (PID2020-113230RB-C21) funded by MICIU/AEI/10.13039/501100011033; Xunta de Galicia (ED431C 2020/11); by Ministry for Digital Transformation and Civil Service and “NextGenerationEU”/PRTR under grant TSI-100925-2023-1; and Centro de Investigación de Galicia “CITIC”, funded by the Xunta de Galicia through the collaboration agreement between the Consellería de Cultura, Educación, Formación Profesional e Universidades and the Galician universities for the reinforcement of the research centres of the Galician University System (CIGUS).

## References

- Hassan Alhuzali and Sophia Ananiadou. 2021. [SpanEmo: Casting multi-label emotion classification as span-prediction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1573–1584, Online. Association for Computational Linguistics.
- Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. 2021. [ViViT: A Video Vision Transformer](#).
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. [wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations](#).
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The Long-Document Transformer](#).
- Ankush Chatterjee, Kedhar Nath Narahari, Meghana Joshi, and Puneet Agrawal. 2019. [SemEval-2019 task 3: EmoContext contextual emotion detection in text](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 39–48, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Ying Chen, Wenjun Hou, Shoushan Li, Caicong Wu, and Xiaoqiang Zhang. 2020. [End-to-End Emotion-Cause Pair Extraction with Graph Convolutional Net-](#)

- work. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 198–207, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised Cross-lingual Representation Learning at Scale](#).
- 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. 2020. [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, Online. Association for Computational Linguistics.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale](#).
- Timothy Dozat and Christopher D Manning. 2017. [Deep Biaffine Attention for Neural Dependency Parsing](#). In *International Conference on Learning Representations*.
- Timothy Dozat and Christopher D. Manning. 2018. [Simpler but more accurate semantic dependency parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 484–490, Melbourne, Australia. Association for Computational Linguistics.
- Soumya Dutta and Sriram Ganapathy. 2024. [HCAM – Hierarchical Cross Attention Model for Multi-modal Emotion Recognition](#).
- 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, Online. Association for Computational Linguistics.
- Mauajama Firdaus, Hardik Chauhan, Asif Ekbal, and Pushpak Bhattacharyya. 2020. [MEISD: A multi-modal multi-label emotion, intensity and sentiment dialogue dataset for emotion recognition and sentiment analysis in conversations](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4441–4453, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Yuan Gong, Yu-An Chung, and James Glass. 2021. [AST: Audio Spectrogram Transformer](#).
- Ian J. Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, Yingbo Zhou, Chetan Ramesh, Fangxiang Feng, Ruifan Li, Xiaojie Wang, Dimitris Athanasakis, John Shawe-Taylor, Maxim Milakov, John Park, Radu Ionescu, Marius Popescu, Cristian Grozea, James Bergstra, Jingjing Xie, Lukasz Romaszko, Bing Xu, Zhang Chuang, and Yoshua Bengio. 2013. [Challenges in Representation Learning: A report on three machine learning contests](#).
- Tao Ji, Yuanbin Wu, and Man Lan. 2019. [Graph-based dependency parsing with graph neural networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2475–2485, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled Weight Decay Regularization](#).
- Fuyan Ma, Bin Sun, and Shutao Li. 2022. [Facial Expression Recognition With Visual Transformers and Attentional Selective Fusion](#). *IEEE Transactions on Affective Computing*, 14(2):1236–1248.
- André Martins, Miguel Almeida, and Noah A. Smith. 2013. [Turning on the turbo: Fast third-order non-projective turbo parsers](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 617–622, Sofia, Bulgaria. Association for Computational Linguistics.
- Ryan McDonald. 2006. [Discriminative training and spanning tree algorithms for dependency parsing](#). *University of Pennsylvania, PhD Thesis*.
- Cam Van Thi Nguyen, Tuan Mai, Son The, Dang Kieu, and Duc-Trong Le. 2023. [Conversation Understanding using Relational Temporal Graph Neural Networks with Auxiliary Cross-Modality Interaction](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15154–15167, Singapore. Association for Computational Linguistics.
- Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. 2022. [Expanding Language-Image Pretrained Models for General Video Recognition](#).

Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. [MELD: A multimodal multi-party dataset for emotion recognition in conversations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 527–536, Florence, Italy. Association for Computational Linguistics.

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust Speech Recognition via Large-Scale Weak Supervision](#).

Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. [You Only Look Once: Unified, Real-Time Object Detection](#).

Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2017. [Modeling Relational Data with Graph Convolutional Networks](#).

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention Is All You Need](#).

Fanfan Wang, Zixiang Ding, Rui Xia, Zhaoyu Li, and Jianfei Yu. 2023a. [Multimodal Emotion-Cause Pair Extraction in Conversations](#). *IEEE Transactions on Affective Computing*, 14(3):1832–1844.

Fanfan Wang, Heqing Ma, Rui Xia, Jianfei Yu, and Erik Cambria. 2024. [Semeval-2024 task 3: Multimodal emotion cause analysis in conversations](#). In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 2022–2033, Mexico City, Mexico. Association for Computational Linguistics.

Wenhai Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, Xiaogang Wang, and Yu Qiao. 2023b. [InternImage: Exploring Large-Scale Vision Foundation Models with Deformable Convolutions](#).

Penghui Wei, Jiahao Zhao, and Wenji Mao. 2020. [Effective inter-clause modeling for end-to-end emotion-cause pair extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3171–3181, Online. 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.

Wenjie Zheng, Jianfei Yu, Rui Xia, and Shijin Wang. 2023. [A facial expression-aware multimodal multi-task learning framework for emotion recognition in multi-party conversations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational*

*Linguistics (Volume 1: Long Papers)*, pages 15445–15459, Toronto, Canada. Association for Computational Linguistics.

## A Appendix

Input	Description
$U_i$	Utterance $i$ , defined as $U_i = (W_i, s_i, \varepsilon_i, \mathbf{X}_i, \mathbf{a}_i)$ .
$W_i$	Word sequence of $U_i$ as $W_i = (w_1, \dots, w_{\ell w, i })$
$s_i$	Speaker of $U_i$ , where $s_i \in \mathcal{S}$
$\varepsilon_i$	Emotion triggered in $U_i$ , where $\varepsilon_i \in \mathcal{E}$
$\mathcal{S}$	Set of speakers in the dataset.
$\mathcal{E}$	Set of annotated emotions.
$\mathbf{X}_i$	Sequence of frames of $U_i$ as $\mathbf{X}_i = (\mathbf{x}_1, \dots, \mathbf{x}_{\ell x, i })$
$\mathbf{x}_k^{(i)}$	Specific frame of $\mathbf{X}_i$ , where $\mathbf{x}_k^{(i)} \in \mathbb{R}^{h \times w \times 3}$ .
$\mathbf{a}_i$	Sampled audio signal of $U_i$ , where $\mathbf{a}_i \in \mathbb{R}^{\ell a, i }$ .
$\ell w, i $	Length of the sequence $W_i$ .
$\ell x, i $	Length of the sequence $\mathbf{X}_i$ .
Encoder	Description
$\mathbf{u}_i$	Encoder hidden representation of $U_i$ from BERT, where $\mathbf{u}_i \in \mathbb{R}^{1024}$ .
$\mathbf{W}_i$	BERT word embeddings of $W_i$ as $\mathbf{W}_i = (\mathbf{u}_i, \mathbf{w}_1^{(i)}, \dots, \mathbf{w}_{\ell w, i }^{(i)})$ .
$\tilde{\mathbf{x}}_i$	Visual hidden representation for $U_i$ , obtained as $\tilde{\mathbf{x}}_i = \text{LSTM}_x^{-1}(\text{ViT}(\mathbf{X}_i)) \in \mathbb{R}^{d_v}$ .
$\tilde{\mathbf{a}}_i$	Audio hidden representation for $U_i$ , obtained as $\tilde{\mathbf{a}}_i = \text{LSTM}_a^{-1}(\text{wav2vec}(\mathbf{a}_i)) \in \mathbb{R}^{d_a}$ .
$\tilde{\mathbf{u}}_i$	Multimodal representation for $U_i$ as $\tilde{\mathbf{u}}_i = (\mathbf{u}_i   \tilde{\mathbf{x}}_i   \tilde{\mathbf{a}}_i)$ .
Decoder	Description
$\Phi$	Arbitrary feed-forward network.
$\mathbf{c}_i$	Cause embedding for $U_i$ as $\mathbf{c}_i = \Phi_c(\mathbf{u}_i) \in \mathbb{R}^{d_G}$ .
$\mathbf{e}_i$	Effect embedding for $U_i$ as $\mathbf{e}_i = \Phi_e(\mathbf{u}_i) \in \mathbb{R}^{d_G}$ .
$\bar{\mathbf{c}}_i$	Emotion cause embedding for $U_i$ as $\bar{\mathbf{c}}_i = \Phi_{c, \varepsilon}(\mathbf{u}_i) \in \mathbb{R}^{d_G}$ .
$\bar{\mathbf{e}}_i$	Emotion effect embedding for $U_i$ as $\bar{\mathbf{e}}_i = \Phi_{e, \varepsilon}(\mathbf{u}_i) \in \mathbb{R}^{d_G}$ .
$\mathbf{C}$	Matrix of cause embeddings as $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_m)$ .
$\mathbf{E}$	Matrix of effect embeddings as $\mathbf{E} = (\mathbf{e}_1, \dots, \mathbf{e}_m)$ .
$\bar{\mathbf{C}}$	Matrix of emotion cause embeddings as $\bar{\mathbf{C}} = (\bar{\mathbf{c}}_1, \dots, \bar{\mathbf{c}}_m)$ .
$\bar{\mathbf{E}}$	Matrix of emotion effect embeddings as $\bar{\mathbf{E}} = (\bar{\mathbf{e}}_1, \dots, \bar{\mathbf{e}}_m)$ .
$\mathcal{W}$	Trainable weights for the first biaffine module, where $\mathcal{W} \in \mathbb{R}^{d_G \times d_G}$ .
$\mathcal{W}_\varepsilon$	Trainable weights for the second biaffine module, where $\mathcal{W}_\varepsilon \in \mathbb{R}^{d_G \times  \mathcal{E}  \times d_G}$ .

Table 2: Symbol notation.