

MIPS at SemEval-2024 Task 3: Multimodal Emotion-Cause Pair Extraction in Conversations with Multimodal Language Models

Zebang Cheng^{1*}, Fuqiang Niu^{1*}, Yuxiang Lin¹
Zhi-Qi Cheng², Bowen Zhang^{1†}, Xiaojiang Peng¹

¹Shenzhen Technology University ²Carnegie Mellon University

zebang.cheng@gmail.com nfq729@gmail.com lin.yuxiang.contact@gmail.com zhiqic@cs.cmu.edu

zhang_bo_wen@foxmail.com pengxiaojiang@sztu.edu.cn

<https://github.com/MIPS-COLT/MER-MCE.git>

Abstract

This paper presents our winning submission to Subtask 2 of SemEval 2024 Task 3 on multimodal emotion cause analysis in conversations. We propose a novel Multimodal Emotion Recognition and Multimodal Emotion Cause Extraction (MER-MCE) framework that integrates text, audio, and visual modalities using specialized emotion encoders. Our approach sets itself apart from top-performing teams by leveraging modality-specific features for enhanced emotion understanding and causality inference. Experimental evaluation demonstrates the advantages of our multimodal approach, with our submission achieving a competitive weighted F1 score of 0.3435, ranking third with a margin of only 0.0339 behind the 1st team and 0.0025 behind the 2nd team.

1 Introduction

Emotion-Cause pair extraction in conversations has garnered significant attention due to its wide-ranging applications, such as optimizing customer service interactions and tailoring content recommendations based on user emotions (Xia and Ding, 2019). However, a fundamental challenge lies in identifying the causal determinants of emotional states. Recent research emphasizes the exploration of causes triggering emotions from multimodal data (Li et al., 2022), followed by further generation tasks based on multimodal emotional cues (Xu et al., 2024).

Practical conversations often exhibit multimodal cues through facial expressions, vocal changes, and textual content. Recognizing the significance of multimodal information, Wang et al. (2023) proposed the task of Multimodal Emotion-Cause Pair Extraction in Conversations (MECPE) as a critical step towards understanding the fundamental elicitors of emotions. The Emotion-Cause-in-

*Equal contributions, collaborated with CMU.

†Corresponding author.

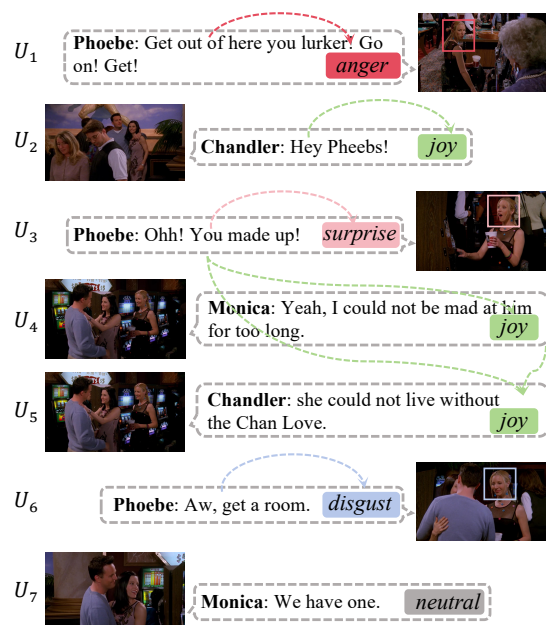


Figure 1: An example of an annotated conversation from the ECF dataset. Dashed lines connect each emotion label to its corresponding cause utterance, illustrating the emotion-cause pairs present in the conversation. The image modality provides additional context and cues for understanding the expressed emotions.

Friends (ECF) dataset, introduced in SemEval 2024 Task 3 (Wang et al., 2024), incorporates additional modalities such as images and audio alongside the original textual data, enabling a more realistic and comprehensive approach to emotion understanding. Figure 1 illustrates an example of an annotated conversation in the ECF dataset, where variations in facial expressions directly mirror the emotions expressed by the characters.

To address the MECPE task, we propose the Multimodal Emotion Recognition-Multimodal Cause Extraction (MER-MCE) framework, building upon the two-step approach introduced by Wang et al. (2023). Our method adopts a two-stage process to predict emotions and identify emotion causes in multimodal conversations. In the first stage, MER-MCE leverages text, audio, and image modalities

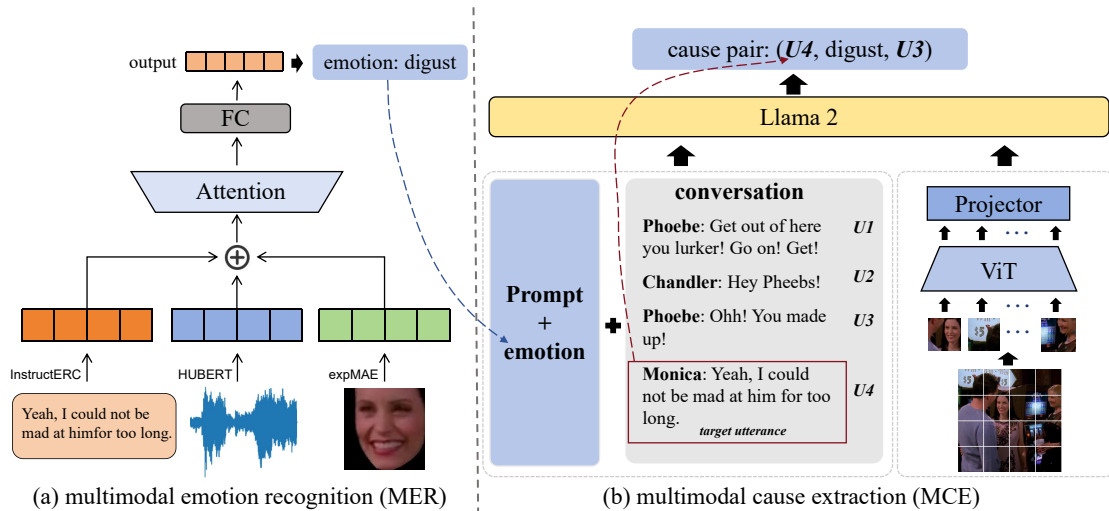


Figure 2: The architecture of our proposed MER-MCE framework for multimodal emotion-cause pair extraction in conversations. The framework consists of two main stages: (a) Multimodal Emotion Recognition (MER), which utilizes specialized emotion encoders to extract modality-specific features from text, audio, and visual data, and (b) Multimodal Cause Extraction (MCE), which employs a Multimodal Language Model to integrate contextual information from the conversation and visual cues to identify the utterances that trigger the recognized emotions.

for emotion prediction, utilizing state-of-the-art models specifically designed for capturing emotional cues from each modality. This approach sets our work apart from the first-place team in the competition, who relied solely on the textual modality for emotion recognition, and the second-place team, who employed a general-purpose model, ImageBind (Girdhar et al., 2023), for extracting visual and audio features.

In the second stage, considering the complexity of analyzing emotion causes for each utterance, we employ a Multimodal Large Language Model (LLM) to dissect the visual and textual modalities and discern the origins of each emotion. By leveraging the power of Multimodal LLMs, our approach can effectively capture the intricate relationships and dependencies present in real-world conversations, enabling a more nuanced and accurate identification of emotion causes.

The MER-MCE framework achieved notable results in Subtask 2 of SemEval 2024 Task 3, ranking third with a weighted F1 score of 0.3435, only 0.0339 behind the first-place team and 0.0025 behind the second-place team. We evaluate the two stages of our model separately to analyze their efficacy and limitations, providing valuable insights into the inherent challenges of the MECPE task. The main contributions of this paper are as follows:

- We propose the MER-MCE framework, a novel two-stage approach for Multimodal Emotion-Cause Pair Extraction in Conversations, leveraging state-of-the-art models for

emotion recognition and Multimodal LLMs for cause extraction.

- We demonstrate the effectiveness of incorporating multiple modalities, including text, audio, and visual information, in both emotion recognition and cause extraction stages, leading to improved performance compared to approaches relying on a single modality or general-purpose feature extractors.
- Through comprehensive evaluation and analysis of the MER-MCE framework on the ECF dataset, we provide valuable insights into the challenges and opportunities in the field of multimodal emotion-cause pair extraction, paving the way for future research and advancements.

2 System Overview

In this work, we propose a Multimodal Emotion Recognition and Multimodal Cause Extraction (MER-MCE) framework for the task of multimodal emotion cause prediction and extraction (MECPE) in conversational settings. As illustrated in Figure 2, our MER-MCE model comprises two key modules: a multimodal emotion recognition (MER) module and a multimodal cause extraction (MCE) module, designed to work in tandem to tackle the intricate challenge of identifying emotions and their underlying causes from multimodal conversational data. Following we describe the structure of the entire system in detail.

2.1 Multimodal Emotion Recognition

Textual Modality. To comprehensively capture the rich semantic and contextual information present in real-world conversational content, our multimodal emotion recognition (MER) module adopts a carefully designed approach that leverages state-of-the-art models tailored for different modalities. For the textual modality, we employ the Instruction-ERC model proposed by [Lei et al. \(2023\)](#), which incorporates a domain demonstration recall module based on semantic similarity to enhance feature extraction. The textual features are extracted in the form of logits, capturing the nuanced semantic representations of the conversational utterances.

Auditory Modality. Recognizing the importance of auditory cues in conveying emotions, we utilize the HuBERT model proposed by [Hsu et al. \(2021\)](#) to process the audio modality and extract hidden states as acoustic features. These acoustic features encapsulate the rich tonal and prosodic information present in the audio signals, complementing the textual and visual modalities.

Visual Modality. For the visual modality, we first employ the OpenFace ([Baltrusaitis et al., 2018](#)) open-source tool to extract facial regions from video clips, allowing our visual model to focus specifically on facial expression recognition. Subsequently, we leverage the expMAE ([Cheng et al., 2023](#)) model to extract both static and dynamic features of facial expressions simultaneously. This dual-feature extraction approach captures the nuanced and time-varying aspects of facial expressions, which are known to be crucial indicators of emotional states.

Multimodal Fusion Mechanism. To effectively fuse the complementary information from these diverse modalities, we employ an attention-based multimodal fusion mechanism, as depicted in [Figure 2\(a\)](#). The input features from each modality are first mapped to a common 128-dimensional space, and ReLU activation and dropout regularization are applied to introduce non-linearity and improve generalization. We then compute attention weights via dot products between the features, allowing the model to dynamically attend to the most relevant and informative cues from each modality. The resulting fused representation captures the synergistic and complementary information across modalities, enabling a comprehensive understanding of the expressed emotions.

2.2 Multimodal Cause Extraction

Building upon the predicted emotions from the multimodal emotion recognition module, we propose a generative approach for multimodal cause extraction (MCE), harnessing the power of Multimodal Language Models (LLMs). As depicted in [Figure 2\(b\)](#), we adopt the MiniGPTv2 ([Chen et al., 2023](#)) model, a state-of-the-art multimodal LLM based on the LLaMA-2 ([Touvron et al., 2023](#)) architecture, as the backbone of our MCE module. This model is designed to integrate both visual and textual information, enabling it to extract emotional causes from multimodal conversational data.

Image Processing. The image processing component of our MCE module employs the Vision Transformer (ViT) ([Dosovitskiy et al., 2020](#)) model, which divides the input image into patches and extracts visual tokens representing these patches. Notably, in contrast to the image preprocessing approach used in the multimodal emotion recognition stage, we feed the complete image to the ViT encoder. This design choice allows our model to capture comprehensive scene information and relationships among individuals, providing a holistic visual context for emotion cause extraction. The extracted visual tokens are then mapped to the textual space using a linear projection layer, facilitating the seamless integration of visual and textual information within the multimodal LLM architecture.

Text Processing. To effectively incorporate contextual information from the conversation, we adopt a prompt-based approach in the textual processing component. The prompt template, illustrated in [Figure 3](#), consists of two key elements: the prompt and the target utterance being queried. The prompt encompasses the conversation content preceding the target utterance, the speaker associated with the queried utterance, and the predicted emotion label obtained from the multimodal emotion recognition module.

Multimodal Cause Extraction. The integration of image and textual data is facilitated by a trainable LLaMA2-chat (7B) model, which processes the multimodal inputs and generates natural language responses to the posed inquiries. These responses are then subjected to a similarity matching process against utterances from the historical conversation dataset. This step allows us to identify the most relevant utterance that potentially triggered the recognized emotion, culminating in the extraction of the ultimate emotion cause utterance. By

seamlessly integrating visual and textual contextual information, our model can capture the intricate relationships and dependencies present in real-world conversations, enabling a more nuanced and flexible representation of emotion causes. This generative approach paves the way for more accurate and interpretable emotion cause extraction, a critical component in the overall task of multimodal emotion cause prediction and extraction.

prompt: Now you are expert of sentiment and emotional analysis. The following conversation noted between [*conversation*] involves several speakers. Please infer, considering the conversation content and the facial expressions of the characters in the image, which utterance leads to [*speaker*] feeling [*emotion*].
target: [*conversation U_i*].

Figure 3: Prompt template for guiding the Multimodal LLM in sentiment analysis and emotion cause extraction from conversational data.

3 Experiments

3.1 Experimental Setup

We evaluate our MER-MCE model on the Emotion-Cause-in-Friends (ECF) dataset (Poria et al., 2018), an extension of the multimodal MELD dataset that includes emotion cause annotations in addition to the original emotion labels. The dataset provides annotations for utterances that trigger the occurrence of emotions, enabling the study of emotion-cause pairs in conversations. For Subtask 2, the labeled data is divided into "train," "dev," and "test" subsets, containing 1001, 112, and 261 conversations, respectively.

Our experimental setup involves extracting features from textual, audio, and visual modalities using various state-of-the-art pretrained models. For the textual modality, we use models such as XLNet (Yang et al., 2019), RoBERTa (Liu et al., 2019), BERT (Devlin et al., 2018), and InstructERC (Lei et al., 2023) to capture semantic and contextual information. For the audio modality, we employ models like VGGish (Hershey et al., 2017), wav2vec (Schneider et al., 2019), and HUBERT (Hsu et al., 2021) to extract features that encapsulate tonal and prosodic information. In the case of the visual modality, we use pretrained visual models, such as MANet (Zhao et al., 2021), ResNet (He et al., 2016), and expMAE (Cheng et al., 2023), to capture nuanced and time-varying aspects of facial expressions conveying emotional information.

3.2 Evaluation Metrics

The primary objective of Subtask 2 is to predict emotion-cause pairs for non-neutral categories based on the provided conversations. The performance of the participating models is evaluated using a weighted average of the F1 scores across six emotion categories: anger, disgust, fear, joy, sadness, and surprise. This weighted average F1 score provides a comprehensive evaluation of the model’s ability to accurately predict emotion-cause pairs while considering the imbalanced nature of the dataset.

Further details on the dataset, experimental setup, and evaluation metrics can be found in the supplementary material.

3.3 Emotion Recognition Analysis

We conducted an extensive experimental evaluation of the Multimodal Emotion Recognition (MER) component within the MER-MCE framework. Table 1 presents the weighted F1 scores for various state-of-the-art models evaluated in the emotion recognition task, leveraging features extracted from textual, audio, and visual modalities.

Our analysis reveals that the textual modality, which captures rich semantic information conveyed through conversations, plays a crucial role in emotion recognition. The inherent ability of the textual modality to encapsulate abstract semantic information facilitates the extraction of emotional features, resulting in higher scores compared to other modalities in unimodal emotion recognition.

To exploit the complementary nature of different modalities, we performed multimodal feature fusion. The experimental results (Table 1) demonstrate that increasing the number of modalities leads to a significant improvement in emotion recognition accuracy, empowering the model to effectively discern emotions within more complex samples. Notably, even features that exhibit relatively lower precision in unimodal emotion recognition contribute positively when integrated into the multimodal fusion framework, highlighting the importance of multimodal approaches in capturing fine-grained emotional nuances.

However, our analysis also uncovers challenges in the visual and audio modalities. In sitcoms containing multiple characters, the OpenFace tool struggles to accurately identify the current speaker, leading to noise in the visual features. Similarly, canned laughter from the audience contributes to

T	A	V	w-avg. F1
XLNet	-	-	0.4418
RoBERTa	-	-	0.5036
BERT	-	-	0.5128
InstructERC	-	-	0.6606
-	VGGish	-	0.2657
-	wav2vec	-	0.4021
-	HUBERT	-	0.4403
-	-	MANet	0.3999
-	-	ResNet	0.4035
-	-	expMAE	0.4104
InstructERC	VGGish	-	0.6729
InstructERC	HUBERT	-	0.6749
InstructERC	-	ResNet	0.6774
InstructERC	-	expMAE	0.6781
-	HUBERT	ResNet	0.5113
-	HUBERT	expMAE	0.5099
InstructERC	VGGish	ResNet	0.6758
InstructERC	VGGish	expMAE	0.6779
InstructERC	HUBERT	ResNet	0.6792
InstructERC	HUBERT	expMAE	0.6807

Table 1: Multimodal emotion recognition results

noise in the audio modality. Consequently, models trained on these modalities exhibit subpar performance compared to the textual modality.

3.4 Cause Extraction Analysis

In the MCE stage, we conducted a comparison of the cause extraction capabilities between different models and the state-of-the-art MECPE-2steps model (Wang et al., 2023), with the test results presented in Table 2. Initially, we employed the same attention model (Lian et al., 2023) structure as in MER for cause extraction. However, this relatively simple model struggled to capture the relationships between utterances. We then explored the ALBEF model (Li et al., 2021) based on the transformer structure, which allowed the model to focus on the connections between different utterances. Nevertheless, limited training data and imbalanced data distribution led to overfitting.

To address these challenges, we transformed the emotion cause extraction task from a traditional discriminative architecture to a generative architecture based on Multimodal LLM, resulting in improved cause extraction accuracy. We utilized the historical conversation window in the Multimodal LLM prompt to retain contextual information within the conversation. Ablation experiments were conducted to investigate the impact of varying numbers of historical conversation windows on cause extraction (Figure 4). To accurately assess the true influence of MCE, we employed the

Method	F1	w-avg. F1
Heuristic	-	0.1864
MECPE-2steps	-	0.3315
Attention	0.3415	0.3403
ALBEF	0.3644	0.3672
MER-MCE(ours)	0.4074	0.4042

Table 2: Multimodal cause extraction results

actual labels from the test dataset for cause extraction instead of relying solely on the emotions predicted in MER. Our experiments revealed that the effectiveness was maximized when the number of historical windows reached 5. However, as the number of windows increased further, the effectiveness gradually decreased due to the complexity of conversations with a larger historical context.




Figure 4: The line graph depicting scores and historical conversation windows.

3.5 Error Analysis of the Entire System

We conducted quantitative and qualitative error analysis on the two stages of our MER-MCE framework to identify key limitations.




Figure 5: The confusion matrix of multimodal emotion recognition result.





Visual Modality	Historical Conversation Content	Label Pair	Pred Pair
	U4: Cat. U5: Yes! You are so smart! I love you. U6: I love you too.	target: U6 emotion: joy cause: U5	target: U6 emotion: joy cause: U5
	U1: I have no idea what you just said. U2: Put Joey on the phone.	target: U2 emotion: anger cause: U1	target: U2 emotion: anger cause: None
	U2: You know what? It really creeps me out ... U3: Sorry. U4: I am so excited!	target: U4 emotion: joy cause: U5	target: U4 emotion: joy cause: U4
	U9: Sure. Okay. U10: Uh , are you crazy? Are you insane? ... U11: Yeah, I ..., I just know it would make me happy.	target: U11 emotion: neutral cause: None	target: U11 emotion: joy cause: U11

Table 3: Analysis of typical predicted emotion-cause pairs generated by the model, with emphasis on samples labeled as 'neutral' that do not have an associated emotion cause.

Analysis of the emotion recognition results using a confusion matrix (Figure 5) revealed that approximately 20% of non-neutral emotion categories were misclassified as neutral, hindering the subsequent cause analysis stage and impacting the overall recall rate. Additionally, class imbalance in the dataset adversely affected the performance of the "disgust" and "fear" categories, which had the lowest number of annotations.

Further analysis of the final Emotion-Cause Pair, based on representative samples in Table 3, highlighted the impact of different scenarios on our model. Facial occlusion in the visual modality (second sample) led to erroneous emotion classification, suggesting the need for more robust visual processing techniques. Strong emotional distractors in the textual modality (fourth sample) misled the model, emphasizing the importance of sophisticated language understanding methods to disambiguate distractors effectively. The real-time setting posed challenges in capturing long-range dependencies and identifying causes in future utterances (third sample), indicating the need for techniques that can model long-range dependencies and incorporate future context.

Despite these challenges, our MER-MCE model demonstrated the ability to accurately predict emotion-cause pairs by leveraging contextual information from both visual and textual modalities (first sample in Table 3). It identified key areas for improvement, including handling facial occlusion, disambiguating emotional distractors, and capturing long-range dependencies in real-time settings.

4 Conclusion

This paper introduces the MER-MCE model for emotion cause analysis in conversations, developed for SemEval 2024 Task 3. Our model leverages multimodal information and Language Models (LLMs) to identify emotion causes in conversational data, considering textual, visual, and audio modalities. MER-MCE achieves a weighted F1 score of 0.3435, ranking third in the competition. The results demonstrate the effectiveness of multimodal approaches in capturing emotional dynamics. Future work will focus on enhancing generalizability and robustness by exploring additional modalities and advanced techniques. We plan to investigate the incorporation of pose estimation, gesture recognition and facial expression analysis to improve the model's ability to detect emotions.

5 Acknowledgements

This work was supported by the National Natural Science Foundation of China (grants 62306184 and 62176165), the Stable Support Projects for Shenzhen Higher Education Institutions (grant 20220718110918001), and the Natural Science Foundation of Top Talent of SZTU (grants GDRC202320 and GDRC202131). Zhi-Qi Cheng acknowledges support from the Air Force Research Laboratory (agreement FA8750-19-2-0200), the Defense Advanced Research Projects Agency (DARPA) grants funded through the GAILA program (award HR00111990063), and the AIDA program (FA8750-18-20018). We also appreciate the organizers of SemEval 2024.

References

- Tadas Baltrušaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. 2018. Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 59–66. IEEE.
- Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. 2023. Minigt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*.
- Zebang Cheng, Yuxiang Lin, Zhaoru Chen, Xiang Li, Shuyi Mao, Fan Zhang, Daijun Ding, Bowen Zhang, and Xiaojiang Peng. 2023. Semi-supervised multimodal emotion recognition with expression mae. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 9436–9440.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Rohit Girdhar, Alaeldin El-Nouby, Zhuang Liu, Manant Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. 2017. Cnn architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*, pages 131–135. IEEE.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Shanglin Lei, Guanting Dong, Xiaoping Wang, Keheng Wang, and Sirui Wang. 2023. Instructerc: Reforming emotion recognition in conversation with a retrieval multi-task llms framework. *arXiv preprint arXiv:2309.11911*.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Wei Li, Yang Li, Vlad Pandealea, Mengshi Ge, Luyao Zhu, and Erik Cambria. 2022. Ecpcc: emotion-cause pair extraction in conversations. *IEEE Transactions on Affective Computing*.
- Zheng Lian, Haiyang Sun, Licai Sun, Jinming Zhao, Ye Liu, Bin Liu, Jiangyan Yi, Meng Wang, Erik Cambria, Guoying Zhao, et al. 2023. Mer 2023: Multi-label learning, modality robustness, and semi-supervised learning. *arXiv preprint arXiv:2304.08981*.
- 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. *arXiv preprint arXiv:1907.11692*.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2018. Meld: A multimodal multi-party dataset for emotion recognition in conversations. *arXiv preprint arXiv:1810.02508*.
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. 2019. wav2vec: Unsupervised pre-training for speech recognition. *arXiv preprint arXiv:1904.05862*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutit Bhoale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Fanfan Wang, Zixiang Ding, Rui Xia, Zhaoyu Li, and Jianfei Yu. 2023. 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.
- 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.

Chao Xu, Yang Liu, Jiazheng Xing, Weida Wang, Mingze Sun, Jun Dan, Tianxin Huang, Siyuan Li, Zhi-Qi Cheng, Ying Tai, et al. 2024. Facechain-imagineid: Freely crafting high-fidelity diverse talking faces from disentangled audio. *arXiv preprint arXiv:2403.01901*.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.

Zengqun Zhao, Qingshan Liu, and Shanmin Wang. 2021. Learning deep global multi-scale and local attention features for facial expression recognition in the wild. *IEEE Transactions on Image Processing*, 30:6544–6556.

A Appendix

A.1 Experimental Data

The ECF dataset, an extension of the multimodal MELD dataset (Poria et al., 2018), includes emotion cause annotations in addition to the original emotion labels. It provides annotations for utterances that trigger the occurrence of emotions, enabling the study of emotion-cause pairs in conversations.

For Subtask 2, the labeled data is divided into "train," "dev," and "test" subsets, containing 1001, 112, and 261 conversations, respectively. We followed this partition to train, validate, and test our MER-MCE model. To evaluate results, we submitted predictions for an additional 655 unlabeled conversations to the CodaLab platform¹.

A.2 Experimental Setup

To evaluate the effectiveness of our MER-MCE model for the MECPE task, we employed various state-of-the-art pretrained models to extract features from the textual, audio, and visual modalities.

For the textual modality, we directly used the conversation text from the annotated files as input to models such as XLNet (Yang et al., 2019), RoBERTa (Liu et al., 2019), BERT (Devlin et al., 2018), and InstructERC (Lei et al., 2023) to extract textual features that capture semantic and contextual information.

For the audio modality, we extracted audio files from the video clips using FFMPEG² and fed them into models like VGGish (Hershey et al.,

2017), wav2vec (Schneider et al., 2019), and HUBERT (Hsu et al., 2021) to obtain audio features that encapsulate tonal and prosodic information.

In the case of the visual modality, we used the OpenFace toolkit to extract facial features from the video clips while masking out the background. We then employed pretrained visual models, such as MANet (Zhao et al., 2021), ResNet (He et al., 2016), and expMAE (Cheng et al., 2023), which were initially trained on facial expression datasets, to extract visual features that capture nuanced and time-varying aspects of facial expressions conveying emotional information.

By leveraging these diverse pretrained models across multiple modalities, we aim to comprehensively capture the rich emotional cues present in the conversations and evaluate the effectiveness of our MER-MCE model in integrating these multimodal features for emotion-cause pair extraction.

A.3 Evaluation Metrics

The primary objective of Subtask 2 is to predict emotion-cause pairs for non-neutral categories based on the provided conversations. Each emotion-cause pair ($p_i = eu_i, ec_i, cu_i$) consists of three essential elements: the index of the emotion utterance eu_i , the emotion category ec_i , and the index of the cause utterance cu_i .

The performance of the participating models is evaluated using a weighted average of the F1 scores across the six emotion categories: anger, disgust, fear, joy, sadness, and surprise. The F1 score for each emotion category j is computed as follows:

$$F_1^j = \frac{2 \times precision^j \times recall^j}{precision^j + recall^j},$$

where $precision^j$ and $recall^j$ are the precision and recall scores for emotion category j , respectively.

The overall performance is determined by the weighted average F1 score, which takes into account the number of samples in each emotion category:

$$w\text{-avg.}F_1 = \frac{\sum_{j=1}^6 n^j \times F_1^j}{\sum_{j=1}^6 n^j},$$

where n^j denotes the number of samples of category j . This weighted average F1 score provides a comprehensive evaluation of the model’s ability to accurately predict emotion-cause pairs across different emotion categories while considering the imbalanced nature of the dataset.

¹<https://codalab.lisn.upsaclay.fr/competitions/16141>

²<https://ffmpeg.org/>