



LLM-Driven Multimodal Opinion Expression Identification

Bonian Jia^{1,*}, Huiyao Chen^{1,*}, Yueheng Sun^{1,†}, Meishan Zhang², Min Zhang²

¹Tianjin University, China

²Institute of Computing and Intelligence, Harbin Institute of Technology (Shenzhen), China

bonianjia@tju.edu.cn, chenhy1018@gmail.com, yhs@tju.edu.cn,
mason.zms@gmail.com, zhangmin2021@hit.edu.cn

Abstract

Opinion Expression Identification (OEI) is essential in NLP for applications ranging from voice assistants to depression diagnosis. This study extends OEI to encompass multimodal inputs, underlining the significance of auditory cues in delivering emotional subtleties beyond the capabilities of text. We introduce a novel multimodal OEI (MOEI) task, integrating text and speech to mirror real-world scenarios. Utilizing CMU MOSEI and IEMOCAP datasets, we construct the CI-MOEI dataset. Additionally, Text-to-Speech (TTS) technology is applied to the MPQA dataset to obtain the CIM-OEI dataset. We design a template for the OEI task to take full advantage of the generative power of large language models (LLMs). Advancing further, we propose an LLM-driven method STOEI, which combines speech and text modal to identify opinion expressions. Our experiments demonstrate that MOEI significantly improves the performance while our method outperforms existing methods by 9.20% and obtains SOTA results.

Index Terms: multimodal opinion expression identification, large language model, speech modal

1. Introduction

Opinion expression identification (OEI) is a critical task of opinion mining in natural language processing (NLP), initially introduced by Breck et al. [1]. It aims to concurrently recognize the text spans that express particular opinions and their sentiment polarity, facilitating a lot of sentiment-related downstream applications such as voice assistants [2], public sentiment monitoring [3, 4] and depression diagnosis [5]. Neural network-based models for OEI have emerged as the predominant approach. Irsoy and Cardie [6] significantly improved the extraction of opinion viewpoints by employing deep bidirectional recurrent neural networks (RNNs), thereby setting new performance standards. Zhang et al. [7] introduced the use of bidirectional long short-term memory (BiLSTM) networks for sentence encoding and incremental decoding processes.

Following the introduction of BERT [8], Xia et al. [9] proposed a novel span-based end-to-end approach, employing BERT for encoding and the BiLSTM model to improve performance. Subsequently, the field of OEI has increasingly been dominated by pre-trained models. Recently, with the emergence of large language models (LLMs) exhibiting superior capabilities, Wu et al. [10] innovated by reformulating the extraction of opinion targets into a question answering (QA) framework. This approach involves the use of LLMs to manage opinion




Figure 1: Opinion expressions in the same sentence show different emotional polarity in different speech scenarios.

mining tasks through the creation of context-specific prompts. Nonetheless, the application of LLMs in OEI remains unexplored. On the other hand, integrating multimodal features, particularly auditory cues, has shown to be advantageous for OEI. Speech conveys more nuanced emotional content than text [11], as the emotional tone can significantly alter the sentiment polarity of the same text, as illustrated in Figure 1. Moreover, inherent speech features like emphasis and pauses are critical in enhancing OEI task accuracy. These auditory characteristics provide supplementary insights absent in text, thereby facilitating a more precise interpretation of the opinion expressions.

In this work, we introduce a multimodal OEI (MOEI) task, focusing on the integration of text and speech to construct authentic OEI scenarios. Acknowledging the inherent challenge presented by real-world speech, which frequently contains interjections and noise, thereby complicating perfect alignment with text, we utilize the open-source, multimodal film and video datasets CMU MOSEI [12] and IEMOCAP [13] to create the CI-MOEI dataset featuring this imperfect alignment. To further enhance our dataset, we apply text-to-speech (TTS) technology to generate synthesized speech for the extensively utilized OEI dataset, MPQA. This synthesized speech integrates as an augmented component of our dataset, culminating in the creation of the CIM-MOEI dataset. It aims to test whether MOEI data from laboratory environments can serve as effective training material to enhance performance.

LLMs have demonstrated significant potential in sentiments analysis [14, 15], motivating us to expand the LLMs-based approach for OEI. First, we design output templates uniquely suited to the OEI task, leveraging the advanced processing power of LLMs. These templates are engineered to facilitate the simultaneous output of all prediction results, thereby enhancing the efficiency. To further refine our approach, we introduce a novel LLM-driven multimodal methodology, designated as STOEI, which combines speech and text modal information to identify opinion expressions. We encode speech into

* Authors contributed equally.

† Corresponding author.

<https://github.com/3cqschr26/LLM-MOEI>




Figure 2: The overall architecture of our method STOEI, where \oplus indicates vectorial concatenation.

speech features, which are then mapped to the textual vector space through an Adapter. By concatenating these speech features with textual vectors, we enable LLMs to effectively process both speech and textual information.

We conduct experiments on these two datasets, and the results demonstrate a significant enhancement in the performance of MOEI when utilizing both text and speech inputs, compared to traditional unimodal (text-only) inputs. Furthermore, our proposed method surpasses existing multimodal techniques by 9.20%, achieving state-of-the-art (SOTA) results. Our code and datasets will be publicly available to facilitate future research.

2. Dataset

In this paper, we propose a novel task named MOEI, focusing here on the case involving two modalities, speech and text, with the aim of combining speech and text information to further improve the performance of the OEI task. Data selection as well as dataset construction is performed to fulfill the task requirements. Next, we will describe these two steps in detail.

2.1. Data Selection

In multimodal research, the conventional assumption of perfect speech-text alignment does not reflect real-world complexities, where issues like vocal interference, audio offset and transcription errors are common. To address this, our study explores the inherent misalignments by utilizing multimodal emotion datasets like CMU MOSEI, featuring manually transcribed YouTube clips, and IEMOCAP, with actor performances that diverge from scripts. These choices highlight the challenges of accurate speech-text synchronization and introduce realistic noise into our analysis. The audio of the above two datasets are derived from real scenes. In addition, we based our speech synthesis on the MPQA dataset, which is commonly used in OEI tasks, and our MOEI annotation guide is inspired by MPQA.

2.2. Dataset Construction

For multimodal emotion datasets, we first establish annotation guidelines aligned with those of the MPQA, aimed at identifying sentiment expressions and their corresponding polarities within sentences. Then, we recruit three university students with majors in English and one native English speaker to serve as annotators. For each piece of data, the three student annotators provide their assessments, and the native English speaker makes the final decision on which annotation to adopt. This process results in an agreement rate of 81% among the student annotators. In instances of uncertainty, the disputed item is referred to an expert in OEI for final determination.

For the OEI dataset MPQA, we select sentences that are labeled with sentiment polarity tags 'POS' or 'NEG' and remove any unspecified content. Subsequently, we employ TTS functionality of Azure¹ as the tool for speech synthesis. To diversify the synthesized voice, five different speakers and ten types of vocal emotions are chosen. These selections are allocated according to the sentiment polarity of the sentences.

3. Method

3.1. LLM-driven OEI

Traditionally, OEI tasks have typically used a sequence labeling approach. Recently, some researchers have begun to explore the use of LLM for emotion recognition in conversation [16]. Nevertheless, this method does not fully exploit the sequence generation potential of LLMs. We advocate for a novel LLM-driven OEI template, which enables the extraction of all opinion expressions along with their respective sentiment polarities.

And then, we redesign the output template of the OEI task. In the original sentence, we add $\langle tag \rangle \langle /tag \rangle$ to mark the emotional opinion words as shown in the figure. Among them, $\langle tag \rangle$ represents the label start, $\langle /tag \rangle$ represents the label end, and the value of the tag is 'pos' or 'neg', which is used to represent the polarity of emotion. An example of the final output is shown on the right side of Figure 2.

3.2. STOEI

It is difficult to identify expressions containing complex emotions based on text alone. In addition, a sentence can also express multiple emotions, and different emotional expressions have different meanings. Speech representations can often convey clues about emotional information. Therefore, we introduce STOEI, a novel LLM-driven method for the MOEI task that integrates four key components: an LLM encoder, a speech encoder, a modality adapter, and an LLM. The comprehensive architecture of our approach is depicted in Figure 2.

Formally, we have a piece of raw speech U , and its corresponding text S . We add a prompt before the text S for task definition. Thus, we obtain a new text $X = (w_1, w_2, \dots, w_n)$, where w denotes a word. Next, we can get the text embedding representation H^{text} of X by using an LLM encoder. Similarly, we input the piece of raw speech U into the speech encoder and extract speech feature embedding A from the final hidden layer.

The core STOEI lies in jointly exploiting text and speech modalities, underpinned by the modality adapter. This integration leverages an architecture inspired by the BLSP frame-

¹<https://learn.microsoft.com/en-us/azure/ai-services/speech-service/text-to-speech>

Table 1: *Statistics of Datasets. #Sent, #POS, and #NEG denote the number of sentences, positive opinion expressions, and negative opinion expressions in the dataset, respectively.*

Dataset	#Sent	#POS	#NEG	Minutes
CI-MOEI	1739	1226	1591	187.7
CIM-MOEI	8592	5241	8800	1273.0
Test	773	454	480	468.4

work [17], comprising three one-dimensional convolutional layers [18] followed by a bottleneck layer [19]. These convolutional layers serve the primary purpose of sub-sampling, which is critical for distilling speech features to an optimal context size conducive to analysis. Following this, the bottleneck layer introduces additional adaptive layers, significantly boosting the model’s ability to both assimilate and interpret nuances from speech data. We use the modality adapter to map speech features to the same vector space as text features, and we get the speech embedding representation H^{speech} . The whole process can be formalized as follows:

$$H^{\text{speech}} = \text{Adapter}(A). \quad (1)$$

Then, we perform a vectorial concatenation operation on the text embedding H^{text} and the speech embedding H^{speech} to obtain a combined representation H that contains multimodal information, which can be represented as:

$$H = H^{\text{text}} \oplus H^{\text{speech}}, \quad (2)$$

where \oplus indicates vectorial concatenation.

After that, we input the combined representation H into the LLM to obtain the output probability distribution O . For training, we use the lora module to decompose the LLM parameter matrix into two low-rank matrices multiplied together, reducing the number of parameters needed for fine-tuning. We employ cross-entropy loss as our training objective:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{c=1}^M \mathbf{y}_{i,t,c} \log(\text{softmax}(O)_{i,t,c}), \quad (3)$$

where $\mathbf{y}_{i,t,c}$ is an one-hot matrix, for sequence i at time step t , it is 1 if the true next word is c , 0 otherwise.

4. Experiments

In this section, we present the dataset construction process, the results compared with other approaches and the related analysis.

4.1. Dataset details

Table 1 reports the statistics of datasets. The CI-MOEI and Test are authentic speech-based MOEI datasets developed through the annotation of the CMU MOSEI and IEMOCAP. The CIM-MOEI refers to an extension of the CI-MOEI, additionally incorporating synthesized speech generated from the MPQA using TTS technology. The CI-OEI and CIM-OEI have been derived by removing the speech components from their respective datasets, retaining only the textual modality data.

4.2. Experiments Setup

Model details. Regarding speech encoder, the Whisper-large-v2 [20] is selected due to its commendable performance in

speech recognition tasks and its capability in emotion recognition. Regarding the large language model, the Vicuna 7B v1.5 [21] is chosen as the experimental base model, and we use the encoder of Vicuna as the text encoder. Vicuna, a conversational assistant, is fine-tuned on shared user dialogues collected from ShareGPT using LLaMA2. The selection of Vicuna is motivated by its enhanced emotional intelligence [22].

Baseline. In the context of speech perception comparative analysis, the LLaMA2 7B [23] is also selected for evaluation. The text-based baseline models subject to comparison include GPT-4 and Bert-BiLSTM-CRF [24]. GPT-4 is the best text-generation LLM. Bert-BiLSTM-CRF is a representative approach for traditional classification models. We also compare an LLM using cross-modal fusion technology, Whispering-LLaMA [25]. It fuses the audio features and the Whisper linear layer into the LLaMA model, injecting it with audio information, while the Whisper linear layer generates key-value pairs in the cross-attention mechanism of the decoder. To better adapt to the needs of MOEI, we replace the base model of Whispering-LLaMA with Vicuna.

Training details. The entirety of our experiments is conducted utilizing an assembly of eight NVIDIA A100 Tensor Core GPUs, with a capacity of 40G. The models is fine-tuned over 15 epochs, employing a learning rate of 5e-5, with a batch size set to 8. The epoch size is determined based on optimal training outcomes. For optimization, the AdamW optimizer is the selected mechanism. About the Low-Rank Adaptation module, the rank of adaptation is established at 8, with the scaling factor for the adaptation set at 16, and a dropout rate of 0.05. By default, the weights of Whisper and LLaMA2 are frozen. The only parts available for training are the modality adapter and the LoRA module.

4.3. Evaluation Metrics

In OEI, model accuracy is determined by an exact match of (start position, end position, sentiment polarity) tuples against gold-standard labels, requiring identical positions and sentiment polarity. Precision (P), recall (R), and the F_1 score are computed to provide a comprehensive evaluation of the model.

4.4. Main Results

Table 2: *The experiment results of traditional OEI methods and speech perception method. The units in the table are %.*

Method	Dataset	P	R	F1
Traditional Text-based OEI				
Bert-BiLSTM-CRF	CI-OEI	49.24	49.45	49.34
GPT4	—	79.74	64.95	71.63
Vicuna	CI-OEI	72.60	73.85	73.22
Text-Speech based MOEI				
Whispering-LLaMA	CI-MOEI	76.14	74.92	75.53
STOEI-Vicuna	CI-MOEI	84.14	85.32	84.73

Speech perception. Analysis of Table 2 shows that the F1 score of GPT4 increases by 22.29% compared to Bert-BiLSTM-CRF when only text input is available, indicating the potential of LLM in OEI. Meanwhile, we find that Vicuna fine-tuned with CI-OEI can exceed the F1 score of GPT4 by 1.59%, which proves that LLM fine-tuning is a promising solution.

Upon combining the auditory modality, both Whispering-LLaMa and STOEI outperform Vicuna, demonstrating that perceiving non-aligned speech indeed aids LLMs in OEI. STOEI surpasses Whispering-LLaMa by a 9.20% margin in terms of the F1 score. These findings highlight the superiority of the modality combination approach employed by our proposed STOEI, underscoring its effectiveness in aiding OEI through the combination of speech modalities.

Table 3: The experimental results on the CIM-MOEI dataset. † indicates that the base model has been fine-tuned on CIM-OEI.

Method	Dataset	P	R	F1
Traditional Text-based OEI				
Bert-BiLSTM-CRF	CIM-OEI	51.09	54.68	52.83
Vicuna	CIM-OEI	76.52	72.67	74.55
Text-Speech based MOEI				
Whispering-LLaMA	CIM-MOEI	76.96	76.63	76.8
STOEI-Vicuna†	CIM-MOEI	86.45	86.17	86.36
STOEI-Vicuna	CIM-MOEI	90.50	89.60	90.09

Further research the impact of speech and text. From Table 3, it becomes clear that an expansion in the training dataset size can lead to further improvements in model performance. For traditional text-based OEI, Bert-BiLSTM-CRF shows a 3.49% increase in the F1 score after fine-tuning on the larger CIM-OEI corpus compared to training on CI-OEI. Similarly, Vicuna exhibits a performance boost of 1.33% when fine-tuned on CIM-OEI versus CI-OEI. Nonetheless, the modest enhancement in the F1 score, despite significant increases in training data volume, indicates that reliance solely on textual input for solving the OEI task has inherent limitations.

Direct fine-tuning of STOEI-Vicuna on CIM-MOEI yields the best performance, surpassing Whispering-LLaMa by 13.29%. This demonstrates the superior capability of STOEI to synergize speech and text features for Opinion Expression Identification compared to other multimodal approaches. Furthermore, a notable improvement of 5.36% is observed compared to the results on the CI-MOEI dataset, indicating that synthesized speech also harbors a wealth of speech features beneficial for the OEI task. Regarding the suboptimal performance of models fine-tuned on text as a basis for STOEI, it is believed that text-based fine-tuning may lead to an over-reliance on the textual modality by the model.

Fine-grain analysis of joint speech-text. In the analysis presented in Fig. 3, the task becomes incrementally challenging for text-only-based LLMs. Specifically, in scenarios of brief sentences with limited contextual breadth, LLMs face difficulties in accurately identifying opinions and discerning their emotional polarities. However, with adequately lengthened contexts, LLMs exhibit improved performance, accurately interpreting the emotional content provided there is sufficient textual information. Nonetheless, when the sentence is too long, text-only-based is not enough for LLMs to perform more accurate OEI. This is because longer sentences contain more complex emotions and semantics, making it more difficult for LLMs to judge. Conversely, the data indicate that both multimodal combination methods yield higher accuracy rates for shorter sentences. This improvement is attributed to the enhanced clarity and prominence of tonal and emotional cues within shorter textual and speech inputs. As sentence length increases, the STOEI model demonstrates superior capability in discerning complex




Figure 3: F1 scores of Whispering-LLaMA, Vicuna(Text-only) and STOEI for different sentence lengths. The methods in the figure were all obtained by training directly on CIM-MOEI, except Vicuna(Text-only), which trained on CIM-OEI.

emotional expressions.

Replacement large language model. As can be seen from the table, the overall trend of the data is similar to that at Vicuna, but the results are not as good as Vicuna. We believe this disparity stems from the higher emotional intelligence of Vicuna, a viewpoint we have corroborated with the research of Wang et al. [22]. And, the trend of the experiment is similar to Vicuna, validating the effectiveness of STOEI.

Table 4: The experiment result of STOEI replacement the LLM. † indicates that the LLM has been fine-tuned on CIM-OEI.

Base Model	Dataset	P	R	F1
Traditional Text-based OEI				
LLaMA2	CI-OEI	67.20	63.24	65.16
LLaMA2	CIM-OEI	71.28	68.38	69.80
Text-Speech based MOEI				
STOEI-LLaMA2	CI-MOEI	0.7209	71.6	71.83
STOEI-LLaMA2†	CIM-MOEI	75.47	72.88	74.15
STOEI-LLaMA2	CIM-MOEI	83.39	80.17	81.75

5. Conclusion

In this work, we established the novel task of multimodal Opinion Expression Identification (MOEI), integrating text and speech to reflect real-world communication nuances. By leveraging open-source datasets CMU MOSEI and IEMOCAP, we created the CI-MOEI dataset, addressing the alignment challenges between speech and text. Additionally, we utilized Text-to-Speech technology on the MPQA dataset to form CIM-OEI, assessing the effectiveness of multimodal data as training material. Employing large language models (LLMs), we developed a novel LLM-driven approach that combines speech and text modalities to help identify opinion expressions. Our experiments demonstrated significant improvements in MOEI performance with this integrated approach, surpassing existing techniques by 9.20% and achieving state-of-the-art (SOTA) results. This study not only highlights the importance of combining textual and auditory information in sentiment analysis but also sets a precedent for future research in leveraging multimodal inputs for deeper emotional and opinion understanding.

Our work has two major limitations. First, we only conducted our experiments on English datasets. While our approach is applicable to other languages, there is a lack of experimental validation in these cases due to the time-consuming and high cost of data annotation. Second, we have only compared a limited number of methods, and do not compare more methods for LLMs because of the large resource consumption.

6. Acknowledgements

This work is supported by the Project of Higher Education Stability Support Program (No.20220618160306001)

7. References

- [1] E. Breck, Y. Choi, and C. Cardie, "Identifying expressions of opinion in context," in *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, M. M. Veloso, Ed., 2007, pp. 2683–2688. [Online]. Available: <http://ijcai.org/Proceedings/07/Papers/431.pdf>
- [2] Y. Ma, Y. Zhang, M. Bachinski, and M. Fjeld, "Emotion-aware voice assistants: Design, implementation, and preliminary insights," in *Proceedings of the Eleventh International Symposium of Chinese CHI*, 2023, pp. 527–532.
- [3] X. Zhang, G. Xu, Y. Sun, M. Zhang, X. Wang, and M. Zhang, "Identifying chinese opinion expressions with extremely-noisy crowdsourcing annotations," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, S. Muresan, P. Nakov, and A. Villavicencio, Eds. Association for Computational Linguistics, 2022, pp. 2801–2813. [Online]. Available: <https://doi.org/10.18653/v1/2022.acl-long.200>
- [4] K. P. Gunasekaran, "Exploring sentiment analysis techniques in natural language processing: A comprehensive review," *CoRR*, vol. abs/2305.14842, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2305.14842>
- [5] M. Sidorov and W. Minker, "Emotion recognition and depression diagnosis by acoustic and visual features: A multimodal approach," in *Proceedings of the 4th International Workshop on Audio/Visual Emotion Challenge*, ser. AVEC '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 81–86. [Online]. Available: <https://doi.org/10.1145/2661806.2661816>
- [6] O. Irsoy and C. Cardie, "Opinion mining with deep recurrent neural networks," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, A. Moschitti, B. Pang, and W. Daelemans, Eds. ACL, 2014, pp. 720–728. [Online]. Available: <https://doi.org/10.3115/v1/d14-1080>
- [7] M. Zhang, Q. Wang, and G. Fu, "End-to-end neural opinion extraction with a transition-based model," *Inf. Syst.*, vol. 80, pp. 56–63, 2019. [Online]. Available: <https://doi.org/10.1016/j.is.2018.09.006>
- [8] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 4171–4186. [Online]. Available: <https://doi.org/10.18653/v1/n19-1423>
- [9] Q. Xia, B. Zhang, R. Wang, Z. Li, Y. Zhang, F. Huang, L. Si, and M. Zhang, "A unified span-based approach for opinion mining with syntactic constituents," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, K. Toutanova, A. Rumshisky, L. Zettlemoyer, D. Hakkani-Tür, I. Beltagy, S. Bethard, R. Cotterell, T. Chakraborty, and Y. Zhou, Eds. Association for Computational Linguistics, 2021, pp. 1795–1804. [Online]. Available: <https://doi.org/10.18653/v1/2021.naacl-main.144>
- [10] L. Wu, Y. Chen, J. Yang, G. Shi, X. Qi, and S. Deng, "Event-centric opinion mining via in-context learning with chatgpt," in *China Conference on Knowledge Graph and Semantic Computing*. Springer, 2023, pp. 83–94.
- [11] P. Wang, S. Zeng, J. Chen, L. Fan, M. Chen, Y. Wu, and X. He, "Leveraging label information for multimodal emotion recognition," *arXiv preprint arXiv:2309.02106*, 2023.
- [12] A. B. Zadeh, P. P. Liang, S. Poria, E. Cambria, and L.-P. Morency, "Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018, pp. 2236–2246.
- [13] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, pp. 335–359, 2008.
- [14] W. Zhang, Y. Deng, B. Liu, S. J. Pan, and L. Bing, "Sentiment analysis in the era of large language models: A reality check," *arXiv preprint arXiv:2305.15005*, 2023.
- [15] X. Sun, X. Li, S. Zhang, S. Wang, F. Wu, J. Li, T. Zhang, and G. Wang, "Sentiment analysis through llm negotiations," *arXiv preprint arXiv:2311.01876*, 2023.
- [16] S. Lei, G. Dong, X. Wang, K. Wang, and S. Wang, "Instructerc: Reforming emotion recognition in conversation with a retrieval multi-task llms framework," *CoRR*, vol. abs/2309.11911, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2309.11911>
- [17] C. Wang, M. Liao, Z. Huang, J. Lu, J. Wu, Y. Liu, C. Zong, and J. Zhang, "Blsp: Bootstrapping language-speech pre-training via behavior alignment of continuation writing," *arXiv preprint arXiv:2309.00916*, 2023.
- [18] A. Hannun, A. Lee, Q. Xu, and R. Collobert, "Sequence-to-sequence speech recognition with time-depth separable convolutions," *arXiv preprint arXiv:1904.02619*, 2019.
- [19] N. Houlsby, A. Giurigu, S. Jastrzebski, B. Morrone, Q. de Larousilhe, A. Gesmundo, M. Attariyan, and S. Gelly, "Parameter-efficient transfer learning for NLP," in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 2019, pp. 2790–2799. [Online]. Available: <http://proceedings.mlr.press/v97/houlsby19a.html>
- [20] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," 2022. [Online]. Available: <https://arxiv.org/abs/2212.04356>
- [21] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez, I. Stoica, and E. P. Xing, "Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality," March 2023. [Online]. Available: <https://lmsys.org/blog/2023-03-30-vicuna/>
- [22] X. Wang, X. Li, Z. Yin, Y. Wu, and J. Liu, "Emotional intelligence of large language models," *Journal of Pacific Rim Psychology*, vol. 17, p. 18344909231213958, 2023.
- [23] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
- [24] X. Zhang, G. Xu, Y. Sun, M. Zhang, X. Wang, and M. Zhang, "Identifying Chinese opinion expressions with extremely-noisy crowdsourcing annotations," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, S. Muresan, P. Nakov, and A. Villavicencio, Eds. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 2801–2813. [Online]. Available: <https://aclanthology.org/2022.acl-long.200>
- [25] S. Radhakrishnan, C.-H. Yang, S. Khan, R. Kumar, N. Kiani, D. Gomez-Cabrero, and J. Tegnér, "Whispering LLaMA: A cross-modal generative error correction framework for speech recognition," in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, H. Bouamor, J. Pino, and K. Bali, Eds. Singapore: Association for Computational Linguistics, Dec. 2023, pp. 10 007–10 016. [Online]. Available: <https://aclanthology.org/2023.emnlp-main.618>