



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e. g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

- This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.
- A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.
- This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.
- The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.
- When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

Speech Emotion Recognition with ASR Integration

Yuanchao Li



Doctor of Philosophy

THE UNIVERSITY OF EDINBURGH

2025

Abstract

Speech Emotion Recognition (SER) plays a pivotal role in understanding human communication, enabling emotionally intelligent systems, and serving as a fundamental component in the development of Artificial General Intelligence (AGI). However, deploying SER in real-world, spontaneous, and low-resource scenarios remains a significant challenge due to the complexity of emotional expression and the limitations of current speech and language technologies. This thesis investigates the integration of Automatic Speech Recognition (ASR) into SER, with the goal of enhancing the robustness, scalability, and practical applicability of emotion recognition from spoken language.

As a starting point, we explore the interplay between ASR and SER by conducting an in-depth analysis of speech foundation models on emotionally expressive speech, from both acoustic and linguistic perspectives. Our findings uncover inherent limitations of these models: while they achieve strong performance in ASR, they often fail to capture paralinguistic cues essential for emotion recognition, and exhibit emotion-dependent biases. Additionally, we examine how linguistic properties, such as part-of-speech distributions, affective word ratings, and utterance length, interact with ASR performance across different emotion categories. Extensive experiments reveal that ASR errors are not random but follow systematic patterns associated with specific word types and emotional expressions across diverse speaking conditions.

To overcome paralinguistic limitations while utilizing the hierarchical encoding capabilities of speech foundation models, we propose a joint training framework that integrates ASR hidden representations and lexical outputs in a hierarchical manner for SER. To mitigate transcription quality issues, we introduce two complementary ASR error correction strategies that require only a small amount of emotional speech data: a Large Language Model (LLM)-based method that refines N-best hypotheses using emotion-specific prompts, and a Sequence-to-Sequence (S2S) model that leverages discrete acoustic units to correct 1-best hypothesis. Both approaches substantially improve the emotional accuracy of ASR transcriptions with low-resource speech data and enhance downstream SER performance.

Beyond transcription correction, this thesis introduces robust acoustic-lexical fusion frameworks to handle emotion mismatch across modalities and reduce ASR error propagation. We conduct a comprehensive investigation into cross-modal incongruity and propose incongruity-aware fusion strategies that dynamically adapt to modality conflicts, extending solutions beyond SER to include sarcasm and humor detection. Building on this, we develop an ASR-

aware, modality-gated fusion mechanism that integrates ASR error correction with dynamic modality selection in a sequential manner. These fusion strategies achieve strong performance on both controlled and real-world datasets, even under partially erroneous conditions (i.e. ASR transcriptions).

Finally, to reduce reliance on costly labeled emotion data, we present a novel semi-supervised learning framework based on multi-view pseudo-labeling. By leveraging both acoustic similarity metrics and LLM-based confidence estimation, this approach selects high-quality unlabeled samples to augment training. The proposed framework not only outperforms traditional pseudo-labeling methods for SER but also demonstrates generalizability to speech-based Alzheimer’s dementia detection.

Lay Summary

Human speech carries more than just words, it reflects our thoughts, intentions, and emotions. The ability to recognize emotions from speech, known as Speech Emotion Recognition (SER), is increasingly important in areas such as mental health monitoring, elderly care, human-computer interaction, and voice-based assistants. A system that can understand how someone feels, simply by listening to them speak, holds great promise for making technology more empathetic and responsive.

However, recognizing emotions in everyday, natural speech is a complex task. In real-life scenarios, speech can vary greatly depending on the speaker, the environment, and the emotional state. This makes it challenging for machines to accurately detect emotions, especially when the speech is spontaneous, noisy, or emotionally intense.

To address this, the research explores how Automatic Speech Recognition (ASR), technology that converts speech into text, can be used to help improve emotion recognition. But there's a catch: while ASR focuses on what is said, it often loses the emotional tone in how it is said. This creates a gap between recognizing content and capturing emotion.

This project develops new methods to combine both the acoustic (audio-based) and linguistic (text-based) features of speech to better detect emotion. It also uses advanced AI tools, such as large language models, to recover emotional information that might be lost or misinterpreted during speech-to-text conversion, especially when the speaker is upset, excited, or speaking in a non-standard way.

Furthermore, the research introduces techniques that help machines learn from unlabeled or uncertain data, using strategies like confidence scoring and self-training. This makes emotion recognition more robust, flexible, and suitable for real-world applications where clean, annotated data is limited.

The findings are not merely theoretical. They have been tested in a variety of scenarios, including practical applications, and have even been extended to the early detection of conditions such as Alzheimer's disease.

Overall, this research aims to make machines more emotionally intelligent by helping them understand not just what people say, but how they feel when they say it. It represents a meaningful step toward more natural and compassionate human-computer interactions.

Acknowledgements

I would like to express my sincere gratitude to my primary supervisor, Dr. Catherine Lai, and my second supervisor, Prof. Peter Bell, for their invaluable mentorship throughout my PhD. Their guidance and expertise have been instrumental in shaping my research at the intersection of SER and ASR. With their support, I was able to pursue my work freely and without pressure.

I am also deeply grateful to Dr. Korin Richmond for his insightful annual reviews and valuable collaboration during my MSc supervision, and to Prof. Simon King and Prof. Carlos Busso for serving as examiners of this thesis and offering their professional feedback.

My heartfelt thanks go to my master's supervisor, Prof. Tatsuya Kawahara, a lifelong 先生 in my life. His continuous support has been invaluable not only in my academic pursuits but also in my personal growth. I am also grateful to Prof. Nigel Ward and Dr. Carlos Ishi for introducing me to the fascinating domain of prosody and dialogue and for guiding my first research project.

Furthermore, I greatly appreciate the warm, collegial, and professional environment at the Centre for Speech Technology Research. I have learned deeply from the weekly reading groups and talks, and I am thankful to all the members of the group for their generosity and kindness.

I would also like to thank my mentors, Dr. Dimitra Emmanouilidou and Dr. Hannes Gamper, at the Microsoft Research Audio and Acoustics Research Group for their support during and beyond my internship. I am grateful to all the collaborators I have worked with throughout this PhD journey.

Foremost, I thank my family for their constant love and support.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

Yuanchao Li

Contents

Abstract	ii
Lay Summary	iv
Acknowledgements	v
Declaration	vi
1 Introduction	1
1.1 Motivation	1
1.2 Research Questions and Objectives	3
1.3 Thesis Outline and Contributions	6
2 Overview of SER and Background of ASR Integration	9
2.1 Introduction	9
2.2 Emotion Theories	10
2.2.1 Categorical Emotion Theory	10
2.2.2 Dimensional Emotion Theory	10
2.2.3 Alternative Theories	11
2.3 SER as A Machine Learning Task	12
2.3.1 Feature Extraction	12
2.3.2 Machine Learning Models	15
2.3.3 Fusion Techniques	19
2.3.4 Speech Foundation Models	21
2.4 Datasets	24
2.5 Evaluation Metrics	28
2.5.1 Classification Tasks	28
2.5.2 Regression Tasks	29
2.6 Prerequisites for Integrating ASR into SER	30
2.6.1 ASR Task Definition	30
2.6.2 Evaluation Metrics	31
2.6.3 Relationship between SER and ASR	33
2.7 Challenges in SER and ASR Integration	35
2.7.1 Limited Availability of Emotional Speech Data	35
2.7.2 Mismatch Between Speech Models and Emotional Speech	35
2.7.3 Challenges with Lexical Information	36

2.7.4	Challenges with Integration of Audio and Text Modalities	37
2.8	Chapter Summary	37
3	Utilizing Self-Supervised Acoustic Representations for SER	38
3.1	Introduction	38
3.2	Exploring wav2vec 2.0 on Emotional Speech	39
3.2.1	Experimental Setup	40
3.2.2	Probing SER Performance	40
3.2.3	Probing Paralinguistic Properties	43
3.2.4	Probing Layer Correlations	46
3.2.5	Probing Hierarchical Properties	47
3.2.6	Probing Emotion Bias	49
3.2.7	Summary	50
3.3	Hierarchical Fusion with Joint ASR Fine-Tuning	50
3.3.1	Experimental Setup	51
3.3.2	Feature Fusion and Emotion Output	52
3.3.3	Results	54
3.3.4	Summary	56
3.4	Chapter Summary	57
4	Utilizing ASR Transcription for SER	58
4.1	Introduction	58
4.2	Exploring ASR Transcription of Emotional Speech	59
4.2.1	Experimental Setup	59
4.2.2	Experimental Results	59
4.2.3	WER on Emotional Speech	59
4.2.4	Exploring ASR Errors Based on Part-of-Speech	60
4.2.5	Exploring ASR Errors Based on Affective Scores	62
4.2.6	Exploring WER with Utterance Length	63
4.2.7	Exploring ASR Performance on Different Emotions	64
4.2.8	Exploring ASR Confidence Scores Across Emotions	65
4.2.9	Summary	67
4.3	Impact of ASR Transcription Quality on SER	67
4.3.1	Experimental Setup	68
4.3.2	Benchmarking SER Across WERs with Text Only	70
4.3.3	Benchmarking SER Across WERs with Bimodal Fusion	72
4.3.4	Summary	74
4.4	Chapter Summary	75
5	Correcting ASR Errors for SER	76

CONTENTS	ix
5.1 Introduction	76
5.1.1 LLM-Based Approaches	76
5.1.2 S2S-Based Approaches	77
5.2 LLM-Based Emotion Recognition on ASR Transcription	78
5.2.1 Methodology	79
5.2.2 Experimental Setup	81
5.2.3 Experiments and Results	82
5.2.4 Summary	87
5.3 S2S-Based ASR Error Correction and Emotion Recognition	87
5.3.1 Proposed Architecture	89
5.3.2 Experimental Setup	90
5.3.3 Experiments and Results	91
5.3.4 Summary	101
5.4 Chapter Summary	101
6 Addressing Cross-Modal Incongruity in SER	103
6.1 Introduction	103
6.2 Exploration of Incongruity in Multimodal Inputs	103
6.2.1 Datasets	104
6.2.2 Analysis of Inter-Modal Incongruity	105
6.2.3 Results and Discussion	106
6.3 Incongruity-Aware Modality-Gated Fusion	107
6.3.1 Experimental Setup	110
6.3.2 Results and Discussion	112
6.3.3 Analysis and Discussion	114
6.3.4 Integration with LLMs	118
6.3.5 Summary	121
6.4 ASR Error-Robust Modality-Gated Fusion	121
6.4.1 Experimental Setup	121
6.4.2 Proposed Framework	122
6.4.3 Results	124
6.4.4 Summary	125
6.5 Chapter Summary	125
7 Semi-Supervised Labeling for Effort-Saving SER	126
7.1 Introduction	126
7.2 Proposed Model	127
7.2.1 Multi-View Pseudo-Labeling	128
7.2.2 Semi-Supervised Training	131
7.3 Experimental Setup	132

CONTENTS	x
7.3.1 Results and Discussion	133
7.4 Chapter Summary	136
8 Conclusions and Outlook	137
8.1 Review of Thesis Achievements	137
8.2 Toward Reliable and Applicable SER	139
Appendices	
A Research Outputs Related to This Thesis	141
Bibliography	147

Introduction

1.1 Motivation

In recent years, the rapid advancement of machine learning in acoustic and language modeling has significantly contributed to the development of conversational agents and voice-driven robotic systems. At the core of these systems lies a dialogue engine, primarily composed of Automatic Speech Recognition (ASR) and Natural Language Processing (NLP) systems, which facilitates grammatical and syntactic analysis of speech to derive semantic meaning and respond to user queries effectively. However, despite their impressive capabilities in processing spoken content, these systems remain fundamentally limited in their ability to interpret non-verbal communication cues.

In other words, while these systems excel at understanding *what is said*, they lack the ability to comprehend *how it is said*, including vocal tones, emotional expressions, underlying psychological states, and personal traits, collectively known as paralinguistic information. Human speech carries rich emotional and social cues that are crucial for natural interactions, and failing to incorporate these aspects restricts the effectiveness of human-computer communication.

To bridge this gap and enable emotionally intelligent interactions, research fields such as Affective Computing, Computational Paralinguistics, and Social Signal Processing have emerged as key areas of study over the past two decades. Among these, Speech Emotion Recognition (SER) has attracted significant attention due to its potential to equip machines with emotional awareness, fostering more natural, engaging, and empathetic human-computer interactions.

SER plays a crucial role in various applications, including:

- Human-computer interaction, where emotionally aware systems can provide more engaging and adaptive responses (Kwon, Chan, Hao, & Lee, 2003; K. Wang, An, Li, Zhang, & Li, 2015).
- Mental health monitoring, assisting in early diagnosis and assessment of psychological conditions through speech analysis (Cummins, Baird, & Schuller, 2018; Elsayed et al., 2022).

- Customer service automation, where emotion-aware AI can enhance user satisfaction by tailoring responses based on emotional cues (W. Han, Jiang, Li, Schuller, & Ruan, 2020; B. Li, Dimitriadis, & Stolcke, 2019).
- Personalized user experiences, enabling virtual assistants, social robots, and entertainment systems to adapt based on a user's emotional state (Lakomkin, Zamani, Weber, Magg, & Wermter, 2018; Y. Li, Ishi, Inoue, Nakamura, & Kawahara, 2019).

By analyzing vocal attributes such as pitch, intensity, rhythm, spectral features, and temporal dynamics, SER systems aim to infer the speaker's emotional state, enabling intuitive and context-aware interactions.

Despite the remarkable progress in SER over the past decade, several critical challenges hinder its transition from controlled laboratory settings to real-world applications:

1. Limited understanding of how emotional speech is represented in cutting-edge speech models

While speech models have undergone significant improvements, particularly with deep learning and transformer-based architectures, their role in emotionally expressive speech remains underexplored. Most speech models are optimized for word recognition accuracy, often disregarding prosodic and paralinguistic features that carry emotional cues. Additionally, emotion-rich speech is often more spontaneous, containing disfluencies, hesitations, and expressive variations that challenge traditional speech systems. Therefore, it is essential to investigate:

- Whether and to what extent speech representations contain emotion-related information.
- How to utilize representations from cutting-edge models.

A deeper understanding of how state-of-the-art speech models process and interact with emotional speech is essential to bridge the gap between speech recognition (i.e., automatic transcription) and emotion recognition, ultimately improving the effectiveness of SER frameworks.

2. Difficult integration of linguistic information from error-prone automatic transcripts

Advances in ASR offer a promising avenue for enhancing SER by leveraging textual and phonetic information extracted from speech. Integrating ASR with SER enables the fusion of acoustic and linguistic features, enriching SER models with linguistic context. This multimodal approach enhances robustness in noisy environments, multilingual scenarios, and spontaneous speech (Y. Li, 2018; Lian et al., 2024; Scotti, Galati, Sbattella, & Tedesco, 2021). Additionally, ASR-generated text facilitates linguistically-informed feature extraction, complementing prosodic and spectral cues.

However, despite these benefits, existing ASR-enhanced SER methods face several challenges, including:

- Error propagation: ASR results are not perfect, and transcription errors can negatively impact emotion recognition (Amiriparian et al., 2021; Munot & Nenkova, 2019).
- Alignment issues: Synchronizing acoustic and textual representations is nontrivial, as speech and text operate on different time scales (W. Wu, Zhang, & Woodland, 2021; H. Xu et al., 2019).
- Computational complexity: Processing both speech and textual features increases computational costs, making real-time SER deployment challenging (Peng, Lu, Pan, & Liu, 2021).

Addressing these issues requires novel fusion strategies that effectively integrate linguistic and acoustic characteristics while minimizing the impact of ASR errors on emotion prediction.

3. High cost of data labeling

Emotion annotation in speech is highly subjective and labor-intensive (Chou et al., 2024). Unlike speech-to-text transcription, which has objective ground truth, emotion perception varies significantly across individuals, leading to inconsistencies in annotation. As a result, multiple annotators are typically required per speech sample to reach a reliable consensus, making the process time-consuming and expensive.

The high cost of data labeling impedes large-scale data collection, slowing down the development and generalization of SER models. To this end, leveraging a small amount of labeled data to guide learning on a larger pool of unlabeled data is required to:

- Reduce reliance on human annotations while maintaining model performance.
- Enhance model generalization by incorporating diverse, real-world speech data.
- Accelerate training and deployment of SER systems with limited labeled datasets.

1.2 Research Questions and Objectives

The above-mentioned technical challenges, including the understanding of ASR models¹ on emotional speech, the fusion of acoustic features and ASR linguistic features, and the high cost of emotion data labeling, need be addressed to bridge the gap between research and practical applications. This research aims to tackle these challenges with the aim of answering the following questions:

1. What is the interplay between ASR models and SER?

1. When this research project began, concepts such as large pre-trained models and speech foundation models had not yet emerged. At that time, the mainstream understanding of the earliest large models, such as wav2vec 2.0, was still as ASR models. Therefore, in this paper, we retain the definition we originally used when starting this work and continue to refer to them as ASR models.

While ASR and SER both rely on speech signals as input, their fundamental objectives and processing mechanisms differ significantly. The acoustic model in ASR typically operates on short temporal windows, focusing on word recognition for accurate transcription. In contrast, SER systems analyze longer temporal contexts to capture emotional patterns embedded in speech dynamics.

Previous studies have highlighted a negative impact of emotional expressiveness on ASR performance² (Fernandez, 2004), as speech affected by strong emotions often exhibits articulatory distortions, prosodic variations, and non-standard intonations. On the other hand, integrating lexical information into SER models has been shown to enhance emotion classification accuracy (C. M. Lee & Narayanan, 2005). However, given that ASR performance deteriorates in emotional speech, the full potential of ASR-enhanced SER, particularly in scenarios involving unlabeled or spontaneous emotional speech, remains underexplored.

Objectives include:

- Analyzing how ASR systems process emotionally expressive speech and quantify their limitations.
- Examining the impact of ASR errors on downstream SER models, particularly in real-world, spontaneous speech settings.

2. How to improve ASR transcription quality of small-scale emotional speech for SER?

ASR error correction has been extensively explored as a post-processing strategy to enhance transcription quality. However, its application in the context of SER remains relatively underexplored. One key reason is the dependency of most ASR error correction methods on large-scale training data, which is often unavailable in emotional speech scenarios (B. Lin & Wang, 2023; Mani, Palaskar, Meripo, Konam, & Metze, 2020; Tanaka, Masumura, Masataki, & Aono, 2018; S. Zhang et al., 2021). Furthermore, emotional speech introduces unique and systematic ASR errors due to variations in prosody, pitch, and articulation, which existing models, typically trained on neutral or read speech, fail to handle effectively.

This challenge is intensified by the scarcity of labeled emotional speech corpora, hindering the adaptation of general-purpose ASR error correction models to emotion-rich, small-scale datasets.

Objectives include:

- Adapting existing ASR error correction methods to low-resource emotional speech domains, with a focus on learning to correct emotion-specific transcription errors.
- Improving SER performance³ by utilizing corrected ASR transcriptions in both text-only and speech-text fusion approaches.

2. In this thesis, ASR performance refers primarily to the word error rate, unless otherwise specified.

3. In this thesis, SER performance refers primarily to the accuracy, unless otherwise specified.

3. Can we develop robust speech-text fusion frameworks for ASR-enhanced SER to mitigate error propagation?

Multimodal approaches that integrate acoustic and lexical features have consistently outperformed unimodal emotion recognition systems (Ezzameli & Mahersia, 2023; Kumar, Kaushik, & Raman, 2021; Yoon, Byun, & Jung, 2018). Two primary fusion strategies: feature-level fusion (early fusion) and decision-level fusion (late fusion), have been widely explored. Recent advancements, such as attention-based intermediate-level fusion techniques, have demonstrated further performance gains. For instance, state-of-the-art fusion techniques employ enhanced Transformer models or hybrid architectures R. Hasan, Nigar, Mamun, and Paul (2025); Y. Wang, Shen, Xu, Li, and Zhao (2021).

Most existing fusion methods, however, assume flawless textual transcriptions, an assumption that fails in real-world applications (where ASR errors are common due to noisy environments, varied accents, and imperfect transcriptions), unlike in controlled lab settings where speech quality is clean or transcription can be manually processed. Such errors cause distortions in meaning, which can mislead emotion classifiers based on linguistic features and degrade SER performance.

Objectives include:

- Designing robust ASR-enhanced SER frameworks that effectively integrate speech and text cues while mitigating the impact of ASR errors.
- Investigating novel fusion techniques, such as error-aware fusion and adaptive attention mechanisms to enhance SER robustness in real-world scenarios.

4. Can effort-saving strategies be developed for scalable SER to reduce annotation costs?

Semi-Supervised Learning (SemiSL), such as self-training (Z. Zhang, Coutinho, Deng, & Schuller, 2014) and co-training (J. Liu, Chen, Bu, You, & Tao, 2007), is a primary approach for SER on limited data. For self-training, the classifier uses its own prediction to teach itself. Whereas co-training tries to exploit the mutual information between two models (views or feature domains), each of which uses its predictions to teach not only itself but also the other.

While existing SemiSL approaches predominantly focus on selecting high-confidence samples through the trained SER model using the labeled data, relatively little work has explored the robustness of the unlabeled data for augmenting the training set. Furthermore, although this issue has been extensively studied in ASR (Wallington, Kershenbaum, Bell, & Klejch, 2021; H. Zhu et al., 2023), it remains underexplored in SER. In a word, previous SemiSL models in SER use predicted labels from unlabeled data without considering their inherent reliability, which could result in error propagation and bias accumulation with model training.

Objectives include:

- Developing novel SemiSL strategies for SER, incorporating both confidence-based and reliability-aware data selection⁴.
- Exploring co-training methodologies that exploit multi-view learning across various feature encoders.
- Reducing dependence on large-scale annotated emotion datasets and leverage ASR transcription, making SER models more cost-efficient and effort-saving.

1.3 Thesis Outline and Contributions

This thesis work, initiated in early 2021, represents one of the first efforts in the exploration, integration, and application of ASR within SER. Given that publicly available emotional speech datasets remain significantly smaller than those for ASR, and considering the aforementioned challenges, practical deployment of SER continues to lag behind ASR despite years of research. Nevertheless, the advent of powerful pre-trained models has significantly improved the potential for utilizing ASR in SER, as the learned representations are now far richer than those produced by traditional ASR systems. At the core of this thesis is the seamless integration of ASR into SER, aiming to bridge the gap between theoretical advancements and real-world applications. This research delves into multiple dimensions, including pre-trained models, Large Language Models (LLMs), ASR error correction, multimodal fusion, and semi-supervised learning. Beyond technical aspects, the study takes a comprehensive approach by examining SER through interdisciplinary perspectives, incorporating insights from acoustics, linguistics, psychology, and cognitive science.

Furthermore, the findings and methodologies presented in this work have been extended to the domain of Alzheimer's dementia detection, demonstrating the broader applicability of our insights in critical speech-related healthcare tasks. This highlights the generalizability of the proposed approaches and underscores their potential impact on real-world speech processing applications.

This chapter serves as an introduction to the thesis. The remainder of this thesis is structured as follows.

Chapter 2 provides an overview of SER and background of integration of ASR in SER, covering fundamental aspects such as feature extraction, machine learning models, commonly used datasets, and persistent challenges in the field.

4. In this context, 'confidence' refers to the model's ability to select data based on a predefined threshold, where data points with a prediction confidence higher than the threshold are selected for training. On the other hand, 'reliability' refers to selecting data based on the consensus or agreement among multiple models, where data points that are consistently selected across models are considered more reliable and chosen for further training.

Chapter 3 first presents a layer-wise analysis of a series of wav2vec 2.0 models, investigating the characteristics of each ASR layer on emotional speech to better understand the strengths and limitations of advanced ASR representations in SER. Subsequently, this chapter introduces a joint ASR-SER training framework with hierarchical fusion, leveraging ASR outputs to improve the robustness of SER systems against ASR errors. Through an exploration of different ASR outputs and fusion strategies, our experiments demonstrate that incorporating both ASR hidden states and text outputs via a hierarchical co-attention fusion method yields the best SER performance, approaching the baseline results obtained using ground-truth transcripts. The research work in this chapter has been presented in the following publications:

- **Li, Y.,** Bell, P., & Lai, C. (2022). Fusing ASR outputs in joint training for speech emotion recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2022)*. (Y. Li, Bell, & Lai, 2022)
- **Li, Y.,** Mohamied, Y., Bell, P., & Lai, C. (2023). Exploration of a self-supervised speech model: A study on emotional corpora. In *2022 IEEE Spoken Language Technology Workshop (SLT 2022)*. (Y. Li, Mohamied, Bell, & Lai, 2023)

In **Chapter 4**, a detailed word-level analysis is performed on ASR transcriptions of emotional speech. The frequencies of various PoS tags, affective scores, and utterance lengths significantly influence ASR performance. Certain word classes are easier to recognize for one ASR model but may be more difficult for another. Furthermore, SER accuracy can deteriorate substantially with high WER. To further investigate this, two benchmarking experiments are conducted to evaluate SER performance (in both classification and regression tasks) across different WER levels in text-only and bimodal fusion (text+speech) settings. The research work in this chapter has been presented in the following publications:

- **Li, Y.,** Zhao, Z., Klejch, O., Bell, P., & Lai, C. (2023). ASR and Emotional Speech: A Word-Level Investigation of the Mutual Impact of Speech and Emotion Recognition. In *Proceedings of Interspeech 2023*. (Y. Li, Zhao, Klejch, Bell, & Lai, 2023)
- **Li, Y.,** Bell, P., & Lai, C. (2024). Speech emotion recognition with ASR transcripts: A comprehensive study on word error rate and fusion techniques. In *2024 IEEE Spoken Language Technology Workshop (SLT 2024)*. (Y. Li, Bell, & Lai, 2024)

Building upon this, **Chapter 5** introduces two innovative methods for joint ASR error correction and emotion recognition: a Sequence-to-Sequence (S2S) model-based approach and an LLM-based method. Their respective advantages are demonstrated, with the S2S model-based approach performing better on 1-best ASR hypotheses, while the LLM-based method proves more effective on N-best ASR hypotheses. The research work in this chapter has been presented in the following publications:

- **Li, Y.,** Chen, P., Bell, P., & Lai, C. (2024). Crossmodal ASR error correction with discrete speech units. In *2024 IEEE Spoken Language Technology Workshop (SLT 2024)*. (Y. Li, Chen, Bell, & Lai, 2024)

- **Li, Y.**, Gong, Y., Yang, C. H. H., Bell, P., & Lai, C. (2024). Revise, Reason, and Recognize: LLM-Based Emotion Recognition via Emotion-Specific Prompts and ASR Error Correction. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2025)*. (Y. Li, Gong, Yang, Bell, & Lai, 2025)

Chapter 6 investigates cross-modal incongruity in emotion recognition, revealing how different modalities (including speech, text, and vision) can convey conflicting affective cues. Furthermore, ASR errors have been found to introduce inconsistencies between the acoustic and linguistic modalities. To address this challenge, an incongruity-aware modality-gated fusion mechanism is proposed, minimizing the negative impact of ASR errors and enhance multimodal SER. The research work in this chapter has been presented in the following publications:

- Wang*, Y., **Li***, Y., Liang, P. P., Morency, L. P., Bell, P., & Lai, C. (2023). Cross-attention is not enough: Incongruity-aware dynamic hierarchical fusion for multimodal affect recognition. *arXiv preprint arXiv:2305.13583*.⁵ (Y. Wang et al., 2023)
- **Li, Y.**, Bell, P., & Lai, C. (2024). Speech emotion recognition with ASR transcripts: A comprehensive study on word error rate and fusion techniques. In *2024 IEEE Spoken Language Technology Workshop (SLT 2024)*. (Y. Li, Bell, & Lai, 2024)

In **Chapter 7**, a semi-supervised learning framework is introduced, leveraging pre-trained speech models and LLMs for automatic labeling. This approach significantly reduces reliance on manual annotation, making SER more scalable and efficient. The framework is also validated on Alzheimer's dementia detection, further confirming its effectiveness. The research work in this chapter has been presented in the following publications:

- **Li, Y.**, Zhang, Z., Han, J., Bell, P., & Lai, C. (2024). Semi-Supervised Cognitive State Classification from Speech with Multi-View Pseudo-Labeling. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2025)*. (Y. Li, Zhang, Han, Bell, & Lai, 2025)

Finally, Chapter 8 concludes the thesis by summarizing key findings and discussing future directions for building more reliable and practical SER systems.

5. This paper is co-first-authored. Yuanchao proposed the research topic, designed and guided the experiments, rewrote the manuscript, and coordinated discussions with collaborators,

Overview of SER and Background of ASR Integration

2.1 Introduction

Speech Emotion Recognition (SER) refers to the task of automatically identifying a speaker's emotional state based on their speech signal. As an interdisciplinary research area, SER lies at the intersection of signal processing (El Ayadi, Kamel, & Karray, 2011; Kwon et al., 2003; K. Wang et al., 2015), machine learning (Fayek, Lech, & Cavedon, 2017; K. Han, Yu, & Tashev, 2014; Schuller & Batliner, 2013), linguistics (Polzehl, Schmitt, Metze, & Wagner, 2011; Triantafyllopoulos et al., 2022), and psychology (Coutinho & Dibben, 2013; Sobin & Alpert, 1999). It plays a vital role in affective computing and has a growing number of applications, including human-computer interaction (Cowie et al., 2001; Devillers, Tahon, Sehili, & Delaborde, 2015), virtual agents and assistants (J. Hu, Huang, Hu, & Xu, 2022; Y. Li, 2018), health monitoring (Cummins et al., 2018; H. Jiang et al., 2017), and customer service (W. Han et al., 2020; B. Li et al., 2019).

In practice, SER systems are typically designed to perform either categorical classification (assigning speech to predefined emotion labels) or dimensional regression (predicting continuous emotional attributes). The choice of approach depends on data availability and application requirements, with some studies investigating hybrid models that combine both perspectives to improve performance (Tompkins, Emmanouilidou, Deshmukh, & Elizalde, 2023). This chapter presents an overview of SER, covering emotion models, task definitions, features, model architectures, datasets, evaluation metrics, related work, and major research challenges. In addition, this chapter presents the background of the integration of ASR into SER, which leads to this thesis work.

2.2 Emotion Theories

A well-defined emotion representation is essential for computational modeling. Two dominant paradigms: **categorical emotion theory** (also known as discrete emotion theory) (Ekman, Dalglish, & Power, 1999) and **dimensional emotion theory** (also referred to as continuous emotion theory) (Russell, 1980), form the foundation for emotion annotation and recognition strategies.

2.2.1 Categorical Emotion Theory

Categorical models, such as Ekman’s basic emotions framework, classify emotions into discrete categories like happiness, sadness, anger, fear, disgust, and surprise. These are commonly used in classification-based SER systems. We employ categorical annotations for datasets like IEMOCAP and RAVDESS in this thesis (datasets will be introduced in Section 2.4).

2.2.2 Dimensional Emotion Theory

Dimensional models represent emotions along continuous scales. There are four dimensions, which capture affective states via:

- **Valence:** positivity or negativity
- **Arousal:** intensity or activation
- **Dominance:** control or submission
- **Social Rejection:** feeling included or isolated

Among the four, the most common is the Valence-Arousal (VA) model, which we use for datasets such as MSP-Podcast and CMU-MOSEI in this thesis (datasets will be introduced in Section 2.4).

Both categorical and dimensional models have been widely adopted in SER research, with the choice often depending on the nature of the available data and the specific requirements of the application (El Ayadi et al., 2011). The categorical approach provides a straightforward method for labeling emotions into discrete classes such as happiness, anger, sadness, or fear. This model is particularly advantageous in applications where distinct emotional states are needed to drive system responses.

In contrast, the dimensional model characterizes emotions along continuous affective dimensions, typically valence (positive–negative), arousal (high–low), and sometimes dominance and social rejection, enabling more nuanced modeling of emotional states (Fontaine, Scherer, Roesch, & Ellsworth, 2007; Russell, 1980). One of the key motivations behind the dimensional approach is its capacity to capture not only broad emotion categories but also subtle

intra-category variations (Acosta & Ward, 2011). For instance, it allows distinctions between “hot anger” (high arousal, low valence, high dominance) and “cold anger” (low arousal, low valence, high dominance), which would otherwise be grouped under a single label “anger” in a categorical framework (Cowie & Cornelius, 2003).

Importantly, these two representations are not mutually exclusive. Emotions defined categorically can often be located within a dimensional affective space; for example, “anger” is typically associated with high arousal and negative valence, whereas “calm” corresponds to low arousal and positive valence (Russell, 1980). Conversely, dimensional data can often be clustered into meaningful emotion categories, suggesting a degree of conceptual overlap between the two frameworks.

This complementarity has inspired the development of hybrid models that aim to integrate categorical and dimensional perspectives (Park et al., 2021; S. Wu, Falk, & Chan, 2011). Such models seek to combine the interpretability and clarity of discrete emotion labels with the expressive flexibility of dimensional features, thereby enhancing the robustness and generalizability of SER systems, particularly in real-world and cross-cultural contexts where emotional expression is often context-dependent (Morgan, 2019).

2.2.3 Alternative Theories

In addition to the widely used categorical and dimensional theories, alternatives have been proposed to bridge the gap between these two conventional approaches. For example, the Hourglass model (Cambria, Livingstone, & Hussain, 2012) is organized around four independent yet interrelated affective dimensions: Pleasantness (ranging from extreme sadness to extreme joy), Attention (from boredom to heightened interest), Sensitivity (from calmness to rage), and Aptitude (from disgust to trust). Each dimension comprises six levels of activation, capturing the intensity of affective states (e.g., contentment → joy → ecstasy within the Pleasantness dimension). This structure preserves the interpretability and categorical clarity of discrete models, while incorporating the continuous gradation of emotional intensity characteristic of dimensional models. Nevertheless, the practical application of such alternative emotion models remains relatively limited, and further empirical validation is required to assess their robustness and utility in real-world scenarios.

2.3 SER as A Machine Learning Task

SER is framed as a supervised machine learning task, detecting emotions conveyed in speech signals. It basically consists of two necessary steps: feature extraction and emotion classification/regression (Y. Li, Zhao, Kawahara, et al., 2019).

2.3.1 Feature Extraction

Features for SER can be divided into handcrafted and model-learned categories. Handcrafted features are manually designed using domain-specific prior knowledge, whereas model-learned features are automatically extracted from speech signals using deep learning models.

Handcrafted Features

Handcrafted features, such as spectral features, prosodic features, format features, and voice quality features, have been intensively explored and used for SER.

- **Spectral features** represent characteristics of the vocal tract. Commonly used spectral features include Mel Frequency Cepstral Coefficients (MFCC) (Daneshfar, Kabudian, & Neekabadi, 2020; K. Wang et al., 2015), mel log-mel spectrograms (Y. Li, Zhao, et al., 2019; Satt, Rozenberg, & Hoory, 2017), mel frequency filterbank energies (Ancilin & Milton, 2021; Busso, Lee, & Narayanan, 2007), and Linear Prediction Cepstral Coefficients (LPCC) (Nwe, Foo, & De Silva, 2003; Pohjalainen, Fabien Ringeval, Zhang, & Schuller, 2016).
- **Prosodic features** capture variations in speech intonation, rhythm, and stress patterns, which are crucial for conveying emotional information. Key prosodic features include pitch (fundamental frequency, F0), energy, and speech duration. Since prosody reflects emotional state and speaker variability, it plays an essential role in SER (Koolagudi & Rao, 2012; Luengo, Navas, Hernández, & Sánchez, 2005; Rao, Koolagudi, & Vempada, 2013).
- **Voice quality features** focus on the vibratory behavior of the vocal cords and the resulting acoustic characteristics. These features help analyze phonation attributes such as breathiness, roughness, and tremor. Commonly used voice quality features include jitter (frequency variations in the voice signal), shimmer (amplitude fluctuations), harmonic-to-noise ratio, and cepstral peak prominence (the periodicity strength of the speech signal), and glottal closure measures (how completely the vocal folds close) (Jacob, 2016; X. Li et al., 2007; Lugger & Yang, 2007).

Given the extensive variety of handcrafted features, researchers have developed standardized acoustic feature sets for SER. Some of the most widely recognized handcrafted feature sets include INTERSPEECH 2009 Emotion Challenge (IS09) (Schuller, Steidl, & Batliner, 2009), INTERSPEECH 2013 Computational Paralinguistics Challenge (ComParE 2013) (Schuller et al., 2013), and Extended Geneva Minimalistic Acoustic Parameter Set (eGeMAPS) (Eyben et al., 2015).

These feature sets form a broad range of acoustic descriptors, such as spectral, prosodic, and voice-quality related features, along with their statistical functionals. They serve as benchmark feature representations for conventional machine learning approaches to SER. In this thesis, we mainly use the following handcrafted paralinguistic features:

- **Energy** is the perceptual correlate of sound intensity and depends on both amplitude and frequency sensitivity of the ear (Moore, 2012).
- **Harmonics-to-noise ratio (HNR)** evaluates the proportion of harmonic components to aperiodic noise in voice signals. It is a key indicator of hoarseness and vocal quality (Yumoto, Gould, & Baer, 1982).
- **Pitch**, associated with the fundamental frequency (F0), reflects the perceived “highness” or “lowness” of a sound (Boersma, 1993).
- **Formants (F1, F2, F3)** are vocal tract resonances that characterize different speech sounds. F1 corresponds to vowel height, F2 to tongue advancement, and F3 to more fine-grained articulatory features (Peterson & Barney, 1952).
- **Alpha ratio** represents the energy balance between low (50–1 000 Hz) and high (1–5 000 Hz) frequency bands, often used to assess vocal brightness (Barsties & Maryn, 2013).
- **Hammarberg Index** compares the strongest energy peak in the 0–2 000 Hz band to the highest above 2 000 Hz, used to assess vocal effort (Hammarberg, Fritzell, Gauffin, Sundberg, & Wedin, 1980).
- **Formant 1, 2, and 3 relative energy** measures the intensity at F1, F2, and F3, providing insight into articulation and voice projection (Stevens, 1998).
- **Spectral slope (0–500 Hz, 500–1 500 Hz)** describes how quickly energy decreases with frequency. Steep slopes typically indicate breathiness, while flatter slopes indicate tension (Titze, 1994).
- **Harmonic differences like H1–H2 and H1–A3** are used to infer vocal fold behavior, such as breathy or tense phonation (Hanson, 1997).
- **Jitter** reflects frequency perturbation from cycle to cycle, revealing irregularity in vocal fold vibration (Ferrand, 2002; Titze, 1994).
- **Shimmer** measures amplitude variation between cycles, and is often used in pathology detection (Baken & Orlikoff, 2000).

In recent years, however, handcrafted features have been gradually replaced by deep learning model-learned approaches which allow automatic representation learning and high-level data abstraction. Nonetheless, it remains uncertain whether the automatic representations learned by the model can entirely replace the handcrafted features, which is an issue that will be addressed in this thesis.

Model-Learned Features

Unlike handcrafted features, direct learning of speech representations through deep learning approaches reduces the requirements for explicit feature engineering. These features provide high-level representations of the speech signal and are typically learned automatically in an end-to-end manner. Deep learning models, such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Transformers are widely used to directly process waveforms to extract meaningful features for SER. More recently, pre-trained speech foundation models, such as wav2vec 2.0, HuBERT, and Whisper, have emerged as powerful tools for learning universal representation of speech signals, benefiting SER. We will introduce the most widely used speech foundation models in Section 2.3.4.

Lexical Features

As lexical features are also employed in this thesis, we provide a brief overview here. In SER, lexical features are extracted from manually transcribed speech or ASR outputs and are designed to capture affective cues embedded in language use. Common lexical features include *n-grams* (e.g., unigrams and bigrams), *Part-of-Speech (POS) tags* (Bandhakavi, Wiratunga, Massie, & Cooper, 2016), *bag-of-words* (Rozgić, Ananthakrishnan, Saleem, & Hansen, 2012), and *Term Frequency-Inverse Document Frequency (TF-IDF) weights* (Chatterjee et al., 2020). These features capture syntactic and statistical patterns that may correlate with emotional content. Additionally, *affective scores* (Warriner, Kuperman, & Brysbaert, 2013), which assign values of arousal, valence, and dominance to each word based on an affective score dictionary, have also been utilized to determine the emotion conveyed in transcribed speech.

With the advent of deep learning, more sophisticated lexical representations such as *word embeddings* (e.g., Word2Vec (Church, 2017), GloVe (Pennington, Socher, & Manning, 2014)) (Rezaie, Parnianifard, Zegarra Rodriguez, Mumtaz, & Wuttisittikulkij, 2022; Tzirakis, Zhang, & Schuller, 2018) and *contextual word embeddings* (e.g., BERT, RoBERTa) (B. Kim & Vossen, 2021; Padi, Singh, & Lal, 2022) have also been used to encode semantic and contextual information in transcriptions. These representations often complement conventional lexical features by capturing nuanced emotional signals that are context-dependent. Overall, the integration of lexical features, either as standalone inputs or in combination with acoustic and prosodic cues, has proven to significantly enhance the performance of SER systems.

2.3.2 Machine Learning Models

SER employs a variety of machine learning models, which are typically categorized into conventional machine learning approaches and Deep Neural Networks (DNNs).

Conventional Machine Learning Models

Conventional machine learning models for SER typically rely on handcrafted speech features, which serve as inputs to classical machine learning algorithms such as Gaussian mixture models (H. Hu, Xu, & Wu, 2007; Lanjewar, Mathurkar, & Patel, 2015), logistic regression (Jacob, 2017; M. J. Kim, Yoo, Kim, & Kim, 2015), support vector machines (H. Hu et al., 2007; Y.-L. Lin & Wei, 2005; P. Shen, Changjun, & Chen, 2011), random forests (Noroozi, Sapiński, Kamińska, & Anbarjafari, 2017; Zheng, Li, Ban, & Liu, 2018), and K-nearest neighbors (Lanjewar et al., 2015; Mote, Sisman, & Busso, 2024; Venkata Subbarao, Terlapu, Geethika, & Harika, 2021).

Conventional machine learning models were widely used in the early stages of SER research due to their simplicity, interpretability, and relatively low computational requirements. However, they often struggle to capture complex patterns in speech data, requiring extensive feature engineering. Since these models are not used in this thesis, we omit their detailed descriptions.

Deep Learning Models

With advancements in deep learning, DNN-based approaches have become the dominant paradigm in SER due to their ability to automatically learn relevant feature representations from speech signals. The models used in this thesis are outlined below.

- **Feed-Forward Neural Networks (FFNNs)** are a fundamental architecture in artificial neural networks, characterized by the unidirectional flow of information from the input layer through one or more hidden layers to the output layer, without any recurrent connections.

Given an input vector $\mathbf{x} \in \mathbb{R}^d$, the transformation at each layer l of an L -layer FFNN is defined as follows:

$$\mathbf{z}^{(l)} = \mathbf{W}^{(l)}\mathbf{a}^{(l-1)} + \mathbf{b}^{(l)}, \quad l = 1, 2, \dots, L \quad (2.1)$$

$$\mathbf{a}^{(l)} = \phi^{(l)}(\mathbf{z}^{(l)}) \quad (2.2)$$

where $\mathbf{a}^{(0)} = \mathbf{x}$ is the input, $\mathbf{W}^{(l)}$ and $\mathbf{b}^{(l)}$ are the weight matrix and bias vector for the l -th layer, and $\phi^{(l)}$ is the activation function (e.g., ReLU, sigmoid, or tanh). The intermediate variable $z^{(l)}$ represents the pre-activation (linear) output at layer l , i.e., the result of the affine transformation before applying the activation function.

For multi-class classification tasks, the output layer uses the *Softmax* activation function to generate a probability distribution over K classes:

$$\hat{y}_i = \frac{e^{z_i^{(L)}}}{\sum_{j=1}^K e^{z_j^{(L)}}}, \quad i = 1, 2, \dots, K \quad (2.3)$$

The network is trained by minimizing the categorical cross-entropy loss:

$$\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y}) = - \sum_{i=1}^K y_i \log(\hat{y}_i) \quad (2.4)$$

where $\mathbf{y} \in \mathbb{R}^K$ is the one-hot encoded true label and $\hat{\mathbf{y}}$ is the predicted probability vector. For regression tasks, the output layer typically does not include an activation function, allowing it to output continuous values directly:

$$\hat{\mathbf{y}} = \mathbf{z}^{(L)} \quad (2.5)$$

The most commonly used loss function for regression is the Mean Squared Error (MSE):

$$\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 \quad (2.6)$$

where \hat{y}_i and y_i denote the predicted and ground truth values, respectively.

The use of FFNNs for SER has been studied at least the last decade (Fayek, Lech, & Cavedon, 2015; Fayek et al., 2017).

- **Convolutional Neural Networks (CNNs)** (LeCun & Bengio, 1995) are a specialized class of DNNs designed to exploit the spatial and/or temporal structure of input data. Unlike standard fully connected layers where each neuron is connected to all neurons in the previous layer via matrix multiplication, CNNs replace this operation with convolution: a localized, parameter-shared linear transformation. This makes CNNs particularly efficient and effective for tasks involving grid-like data structures, such as images or spectrograms.

A typical CNN is composed of multiple *convolutional blocks*, where each block generally performs a sequence of operations: convolution, non-linear activation, and pooling. In the convolutional layer, instead of a full weight matrix, a set of learnable kernels (also called filters) \mathbf{K} is used. Each kernel has a small receptive field (e.g., 3×3 or 5×5), which significantly reduces the number of parameters and computational cost. The number of output channels (or feature maps) produced by a convolutional layer is denoted as $C^{(l)}$, where l is the layer index.

The output of a convolutional layer can be formally expressed as:

$$h_j^{(l)} = \sigma \left(\sum_{i=1}^{C^{(l-1)}} h_i^{(l-1)} * K_{ij}^{(l)} + b_j^{(l)} \right), \quad \forall j \in \{1, 2, \dots, C^{(l)}\} \quad (2.7)$$

where $*$ denotes the 2D convolution operation, $h_i^{(l-1)}$ is the i -th input channel of the previous layer, $K_{ij}^{(l)}$ is the kernel connecting input channel i to output channel j , $b_j^{(l)}$ is the bias term, and $\sigma(\cdot)$ is a non-linear activation function such as ReLU.

Following the convolution and activation steps, a *pooling layer* is often applied to down-sample the feature maps. Pooling is a form of non-linear aggregation that improves translation invariance and reduces spatial dimensions. The most common types of pooling are max pooling and average pooling (Zhou & Chellappa, 1988), defined over local regions. Max pooling selects the maximum value, whereas average pooling computes the mean or weighted average of the values within a pooling window.

The hierarchical structure of CNNs allows the network to learn increasingly abstract features at deeper layers. The use of CNNs for SER have proven useful (Huang, Dong, Mao, & Zhan, 2014; Y. Li, Zhao, et al., 2019; Peng et al., 2021).

- **Recurrent Neural Networks (RNNs)** (Hopfield, 1982; Rumelhart, Hinton, & Williams, 1986) are designed to process sequential data by allowing information to be propagated through time. RNNs have connections that allow information to persist from one time step to another, making them ideal for tasks where the temporal order of data is crucial, such as time series forecasting, language modeling, and speech recognition. At each time step t , an RNN updates its hidden state \mathbf{h}_t based on the input \mathbf{x}_t and the previous hidden state \mathbf{h}_{t-1} :

$$\mathbf{h}_t = \sigma(\mathbf{W}_h \mathbf{h}_{t-1} + \mathbf{W}_x \mathbf{x}_t + \mathbf{b}_h) \quad (2.8)$$

where $\sigma(\cdot)$ is a non-linear activation function (often tanh or ReLU), \mathbf{W}_h and \mathbf{W}_x are learnable weight matrices, and \mathbf{b}_h is a bias term. The output of the network at time t is typically a function of the hidden state, $\mathbf{y}_t = \mathbf{W}_y \mathbf{h}_t + \mathbf{b}_y$, where \mathbf{W}_y and \mathbf{b}_y are the weight matrix and bias of the output layer.

However, standard RNNs suffer from the vanishing gradient problem, which makes them difficult to train over long sequences due to the diminishing impact of earlier time steps during backpropagation. This issue is mitigated by the Long Short-Term Memory (LSTM) network (Hochreiter & Schmidhuber, 1997), a specialized type of RNN designed to capture long-range dependencies in sequential data.

LSTMs address the vanishing gradient problem by introducing a more sophisticated architecture that uses gates to control the flow of information. Specifically, LSTM units consist of three gates: the forget gate (f_t), the input gate (i_t), and the output gate (o_t). These gates regulate the information retained or discarded at each time step.

The core update equations for an LSTM unit at time step t are as follows:

$$f_t = \sigma(\mathbf{W}_f \mathbf{h}_{t-1} + \mathbf{W}_x \mathbf{x}_t + \mathbf{b}_f) \quad (2.9)$$

$$i_t = \sigma(\mathbf{W}_i \mathbf{h}_{t-1} + \mathbf{W}_x \mathbf{x}_t + \mathbf{b}_i) \quad (2.10)$$

$$o_t = \sigma(\mathbf{W}_o \mathbf{h}_{t-1} + \mathbf{W}_x \mathbf{x}_t + \mathbf{b}_o) \quad (2.11)$$

$$\mathbf{C}_t = f_t \odot \mathbf{C}_{t-1} + i_t \odot \tanh(\mathbf{W}_C \mathbf{h}_{t-1} + \mathbf{W}_x \mathbf{x}_t + \mathbf{b}_C) \quad (2.12)$$

$$\mathbf{h}_t = o_t \odot \tanh(\mathbf{C}_t) \quad (2.13)$$

In these equations, f_t is the forget gate, determining how much of the previous memory \mathbf{C}_{t-1} should be forgotten. i_t is the input gate, controlling how much new information should be added to the cell state. o_t is the output gate, dictating the output of the LSTM unit. \mathbf{C}_t represents the cell state, which carries the long-term memory of the network. \odot denotes element-wise multiplication, and $\tanh(\cdot)$ is the hyperbolic tangent activation function.

The LSTM architecture allows the model to learn long-range dependencies effectively, as the cell state \mathbf{C}_t is capable of preserving information over many time steps without suffering from vanishing gradients.

RNNs have proven to be effective and complementary to CNNs in SER (J. Lee & Tashev, 2015; Y. Li, Zhao, et al., 2019; Mirsamadi, Barsoum, & Zhang, 2017).

- **Transformer** networks (Vaswani et al., 2017) represent a paradigm shift in the field of sequence modeling. Unlike RNNs and LSTM models, which process sequential data in an inherently temporal and recursive manner, Transformers rely entirely on self-attention mechanisms and dispense with recurrence altogether. This architecture enables efficient parallel computation and superior performance on a wide range of sequence tasks, including machine translation, language modeling, and speech recognition.

The core innovation of the Transformer is the self-attention mechanism, which allows each position in the input sequence to attend to all other positions directly. Given an input sequence $\mathbf{X} \in \mathbb{R}^{T \times d_{\text{model}}}$, the self-attention mechanism computes three matrices: queries \mathbf{Q} , keys \mathbf{K} , and values \mathbf{V} , via learned linear transformations:

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_Q, \quad \mathbf{K} = \mathbf{X}\mathbf{W}_K, \quad \mathbf{V} = \mathbf{X}\mathbf{W}_V \quad (2.14)$$

where $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{d_{\text{model}} \times d_k}$ are learnable weight matrices. The output of scaled dot-product attention is computed as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right) \mathbf{V} \quad (2.15)$$

This mechanism enables the model to dynamically weigh the relevance of each position in the sequence for the current computation.

To capture multiple types of relationships in parallel, the Transformer employs *multi-head attention*, which consists of several self-attention operations run in parallel, each with its own set of parameters. The outputs of the heads are then concatenated and linearly transformed:

$$\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Concat}(\text{head}_1, \dots, \text{head}_h) \mathbf{W}^O \quad (2.16)$$

Each attention head is defined as:

$$\text{head}_i = \text{Attention}(\mathbf{QW}_i^Q, \mathbf{KW}_i^K, \mathbf{VW}_i^V) \quad (2.17)$$

Since the Transformer does not have any recurrence or convolution, it requires *positional encodings* to incorporate the order of the sequence. These are summed to the input embeddings and can be either learned or computed using sinusoidal functions.

The original Transformer is composed of an encoder-decoder structure. The encoder consists of a stack of N identical layers, each comprising two sub-layers: a multi-head self-attention mechanism and a position-wise FFNN. Each sub-layer is followed by residual connections (K. He, Zhang, Ren, & Sun, 2016) and layer normalization (Ba, Kiros, & Hinton, 2016).

The decoder also consists of N layers and includes a third sub-layer that performs multi-head attention over the encoder's output. The decoder is autoregressive, meaning it generates tokens one at a time and uses previously generated tokens as additional input.

Transformers have shown substantial superiority in SER compared to CNNs and RNNs (Gao, Wang, Liu, Dang, & Okada, 2023; X. Wang et al., 2021; Y. Wang et al., 2021).

2.3.3 Fusion Techniques

With the integration of speech features, language features, and machine learning models, SER has greatly benefited from multimodal fusion techniques. These techniques can be broadly categorized into feature-level fusion, intermediate-level fusion, and decision-level fusion, although many approaches inherently involve elements of more than one fusion strategy.

Feature-Level and Decision-Level Fusion

Feature-level fusion, also known as early fusion, involves concatenating or combining features extracted from different modalities (e.g., prosodic, spectral, and linguistic features) into a single feature vector before inputting them into a classification model. This approach enables the model to learn joint representations and modality correlations at an early stage. In contrast, decision-level fusion, also referred to as late fusion, processes each modality independently to produce decisions (e.g., predicted emotion labels or confidence scores), which are subsequently combined using ensemble methods such as majority voting, weighted averaging, or rule-based strategies.

For example, Sebastian, Pierucci, et al. (2019) conducted a comprehensive analysis comparing the emotion classification accuracy and F1 score of a CNN-based audio model, an LSTM-based text model, an early fusion model, and three late fusion models. Specifically, the three late fusion methods employed were sum combination, weighted sum combination, and product fusion. Pepino, Riera, Ferrer, and Gravano (2020) explored three training strategies, cold-start, pre-trained, and warm-start, for early and late fusion models. The cold-start strategy trained the fusion model from scratch. In the pre-trained strategy, speech and text models were trained separately, and their weights were used to initialize the corresponding layers in the fusion model. The warm-start approach initialized all layers as in the pre-trained strategy, but trained the entire model, not just the post-fusion layers. Additionally, early and late fusion can be combined to complement each other (Chuang & Wu, 2004).

However, feature-level fusion methods often struggle when input features from multiple modalities differ in temporal characteristics due to strict synchrony requirements. Furthermore, they are limited in modeling intra-modality dynamics effectively. On the other hand, decision-level fusion cannot capture inter-modality dynamics because the dependencies among modalities are not learned (Y. Li, Zhao, & Shen, 2020). To address these limitations, intermediate-level fusion has become the predominant approach in deep learning-based SER.

Intermediate-Level Fusion

Intermediate-level fusion, also known as representation-level or hidden-level fusion, integrates modalities at latent or hidden representation layers, typically within a deep neural network. Instead of directly combining features, each modality is first encoded independently, and their resulting embeddings are fused at one or more intermediate layers of the network.

This strategy allows for more flexible and hierarchical modeling of modality-specific patterns. For instance, (Zadeh, Zellers, Pincus, & Morency, 2016) introduced the memory fusion network, which captures both unimodal and cross-modal interactions through a gated memory mechanism. With the rising popularity of attention mechanisms (Vaswani et al., 2017), attention-based fusion has become a key trend at the intermediate level. Such methods can dynamically

model inter-modality interactions and capture dependencies between modalities. H. Xu et al. (2019) used bi-directional LSTMs to extract speech and text features, using text representations as the “query” to perform attention-based fusion on speech features. Similarly, Yoon et al. (2018) used attention with speech as the “query”. Building on this, the authors later proposed a multi-hop attention mechanism that automatically infers inter-modal correlations (Yoon, Byun, Dey, & Jung, 2019). In their model, relevant textual segments corresponding to the speech signal were identified via multi-hop attention, and these segments were then used to attend to specific parts of the speech signal.

Although intermediate-level fusion often outperforms feature-level and decision-level fusion, it typically requires careful architectural design and greater computational cost.

2.3.4 Speech Foundation Models

With advancements in deep learning algorithms and computational resources, pre-trained speech models have emerged as powerful tools for learning universal representations of speech signals. These models, trained on large-scale speech corpora, composing the functions of both a feature extractor and a deep learning model for SER. Speech foundation models usually consist of multiple layers of Transformers, reducing the requirements of manually stacking neural networks such as CNNs. To this end, conducting SER using speech foundation models usually requires only simple FFNNs for final prediction output. Below are the speech foundation models that used in this thesis.

Wav2vec 2.0

Wav2vec 2.0 (Baevski, Zhou, Mohamed, & Auli, 2020) is a self-supervised learning model designed to learn contextualized representations from speech waveforms. As shown in Figure 2.1, the model architecture consists of three main components: a CNN-based local encoder that extracts a sequence of embeddings from audio as latent representation; a Transformer network for context representation; and a shared quantizer for discretization across multiple languages.

Given an audio input $\mathbf{x} \in \mathbb{R}^T$, the CNN-based local encoder maps the waveform into a sequence of latent feature representations $\mathbf{z} = (z_1, z_2, \dots, z_{T'})$. A context network, based on a Transformer, processes these latent features to produce contextualized embeddings $\mathbf{c} = (c_1, c_2, \dots, c_{T'})$.

During pretraining, a subset of latent feature representations \mathbf{z} is masked, and the model is trained to correctly identify the true quantized targets $q(\mathbf{z})$ among a set of distractors. The objective is a contrastive loss, defined as:

$$\mathcal{L}_{\text{contrastive}} = - \sum_{t \in \mathcal{M}} \log \frac{\exp(\text{sim}(c_t, q(z_t)) / \kappa)}{\sum_{\tilde{q} \in \mathcal{Q}_t} \exp(\text{sim}(c_t, \tilde{q}) / \kappa)} \quad (2.18)$$

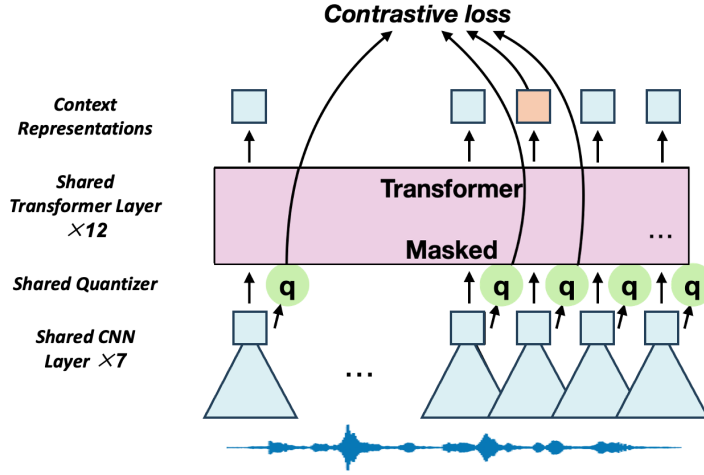


Figure 2.1: Wav2vec 2.0 model.

where \mathcal{M} is the set of masked time steps, $\text{sim}(\cdot, \cdot)$ denotes the cosine similarity, \mathcal{Q}_t includes the true target and negative samples at time t , and κ is a temperature hyperparameter.

Wav2vec 2.0 significantly reduces the amount of labeled data required for competitive performance and has set new benchmarks for semi-supervised and low-resource speech tasks. Due to its ability to encode rich speech features, wav2vec 2.0 has been extensively utilized in SER (L.-W. Chen & Rudnicky, 2023; Gao, Chu, & Kawahara, 2023; Pepino, Riera, & Ferrer, 2021)

HuBERT

HuBERT (Hsu et al., 2021) introduces a self-supervised learning framework for speech representation learning based on masked prediction of discrete hidden units. Inspired by the masked language modeling objective in BERT, HuBERT masks portions of the input speech features and trains the model to predict their corresponding cluster assignments, which serve as pseudo-labels generated through offline clustering.

Same as wav2vec 2.0, the model extracts latent representations \mathbf{z} from waveform \mathbf{x} via a convolutional feature encoder. These features are then processed by a Transformer to produce contextualized embeddings \mathbf{c} . Initially, discrete targets are obtained by applying k-means clustering over the extracted features. During training, the model is optimized with a cross-entropy loss over the masked positions:

$$\mathcal{L}_{\text{CE}} = - \sum_{t \in \mathcal{M}} \log p_{\theta}(q(z_t) | c_t) \quad (2.19)$$

where \mathcal{M} denotes the set of masked time steps, p_{θ} represents the predicted probability distribution over cluster IDs, and $q(z_t)$ is the pseudo-label corresponding to frame t .

The overall architecture is similar to wav2vec 2.0. However, unlike wav2vec 2.0, which directly learns to predict quantized latent representations through a contrastive loss, HuBERT relies on a multi-stage training procedure. After the first training phase, hidden representations from the Transformer are reclustered to produce new targets, enabling iterative refinement. This design allows HuBERT to progressively improve the quality of learned representations without requiring a contrastive objective. As a result, HuBERT often outperforms wav2vec 2.0 across various speech processing applications, including SER (Pastor, Ribas, Ortega, Miguel, & Lleida, 2022; Saliba, Li, Sanabria, & Lai, 2024).

WavLM

WavLM (S. Chen et al., 2022) extends the line of self-supervised speech representation learning by integrating masked prediction with a denoising pre-training task, aiming to learn universal speech representations that are effective across both low-level and high-level speech processing tasks.

Building upon the masked prediction strategy of HuBERT, WavLM introduces additional noise robustness by augmenting the input signals with speech-like disturbances, such as overlapped speech and additive background noise, during pre-training. The model is then trained to predict the correct masked units despite the presence of noise, encouraging the learning of more generalizable and noise-invariant features.

The general architecture and training loss are the same as HuBERT. However, WavLM further incorporates a gated relative position bias mechanism within the Transformer layers, enhancing its capacity to model long-range dependencies and speaker overlaps, which is crucial for tasks such as speaker diarization, separation, and speech recognition in noisy environments.

Thanks to these innovations, WavLM achieves state-of-the-art performance across a wide range of speech tasks, and serves as a highly effective pre-trained model for SER (Diatlova, Udalov, Shutov, & Spirin, 2024; Upadhyay, Busso, & Lee, 2024; J. Yang et al., 2024).

Whisper

Whisper (Radford et al., 2023) is a fully supervised, end-to-end ASR system trained on a large-scale, weakly labeled dataset comprising approximately 6.8×10^5 hours of multilingual and multitask speech data collected from the web. Unlike previous self-supervised models such as wav2vec 2.0, HuBERT, and WavLM, Whisper does not rely on unsupervised pre-training, but instead utilizes massive supervised training to achieve robustness across diverse domains, languages, and acoustic conditions.

The model architecture follows a standard encoder-decoder Transformer framework. The encoder processes log-mel spectrograms $\mathbf{x} \in \mathbb{R}^{F \times T}$ into a sequence of hidden states, where F denotes the number of frequency bins and T is the number of time frames. The decoder is autoregressive and predicts the corresponding text tokens sequentially, conditioned on the encoder outputs.

Formally, given an input spectrogram \mathbf{x} , the model estimates the probability of the output token sequence $\mathbf{y} = (y_1, y_2, \dots, y_N)$ as:

$$P(\mathbf{y} | \mathbf{x}) = \prod_{n=1}^N P(y_n | y_{<n}, \mathbf{x}) \quad (2.20)$$

where $y_{<n}$ denotes the previous tokens, and each conditional probability is computed via the decoder network.

Whisper is trained with a standard cross-entropy loss between the predicted and ground-truth text tokens. Additionally, the model is designed for multitask learning: it can perform language identification, phrase-level timestamps prediction, and multilingual speech recognition within the same framework. Thanks to its scale, diversity of training data, and multitask formulation, Whisper demonstrates strong zero-shot transfer capabilities, achieving competitive or superior performance on many publicly available ASR benchmarks without additional fine-tuning and also demonstrate significant SER performance (Chou, 2024; Fukuda, Kano, Ando, & Ogawa, 2025; Goron, Asai, Rut, & Dinov, 2024).

A brief comparison of the models above is presented in Table 2.1.

2.4 Datasets

High-quality datasets serve as the cornerstone for SER research. This section provides an overview of several widely used SER datasets, with a particular focus on those utilized in this thesis. These datasets play a crucial role in benchmarking SER models and assessing their performance across different experimental settings.

IEMOCAP

The *Interactive Emotional Dyadic Motion Capture* (IEMOCAP) dataset (Busso et al., 2008) is one of the most extensively used English-language corpora for SER. It contains approximately 12 hours of multimodal data, including speech recordings, textual transcriptions, and facial motion capture. The dataset consists of five dyadic conversational sessions, each featuring a male-female pair of professional actors, resulting in a total of ten speakers. Within each session, the participants engage in both pre-scripted dialogues and improvised interactions based on predefined emotional scenarios.

Table 2.1: Comparison of Wav2vec 2.0, HuBERT, WavLM, and Whisper.

Aspect	Wav2vec 2.0	HuBERT	WavLM	Whisper
Training	Self-sup. (contrastive)	Self-sup. (masked pred.)	Self-sup. (denoising)	Supervised
Architecture	Encoder (Transformer)	Encoder (Transformer)	Encoder (Relative bias)	Encoder-Decoder
Input	Waveform	Waveform	Waveform	Log-mel spec
Target	Quantized feats.	K-means labels	K-means + noise	Text tokens
Noise Robustness	No	No	Yes	Yes
Multilingual	No	No	No	Yes
Zero-shot	Limited	Limited	Moderate	Strong

IEMOCAP comprises 151 dialogues, segmented into 10039 utterances, with an average duration of 4.5 s per utterance. Each utterance has been manually annotated by three raters for categorical emotions, including *neutral*, *happy*, *sad*, *angry*, *excited*, *disappointed*, *fearful*, *surprised*, and *frustrated*. Additionally, a dimensional annotation scheme is employed, where utterances are rated on a five-point Likert scale for *valence* (1 = negative, 5 = positive), *activation* (1 = calm, 5 = excited), and *dominance* (1 = weak, 5 = strong). If an utterance conveys mixed emotions, multiple labels may be assigned. Final categorical labels are determined by majority voting, while dimensional scores are obtained by averaging the raters' annotations.

RAVDESS

The *Ryerson Audio-Visual Database of Emotional Speech and Song* (RAVDESS) (Livingstone & Russo, 2018) is a multimodal dataset designed for SER and affective computing research. It contains 7,356 recordings of both spoken and sung expressions, performed by 24 professional actors (12 male, 12 female). The speech component features 16 distinct statements, each delivered in eight emotional states: *neutral*, *calm*, *happy*, *sad*, *angry*, *fearful*, *disgusted*, and *surprised*. All recordings are available in both audio-only and audio-visual formats, enabling cross-modal emotion analysis.

However, as the emotional expressions are acted in this dataset, they may not fully capture the complexity and spontaneity of real-world emotional speech. Nevertheless, RAVDESS remains a valuable benchmark for SER models, particularly when controlled and high-fidelity emotional speech data are required.

MSP-Podcast

The *MSP-Podcast* corpus (Lotfian & Busso, 2017) is a large-scale, naturalistic English emotional speech dataset constructed from publicly available podcast recordings. Developed by the Multimodal Signal Processing (MSP) Lab at the University of Texas at Dallas, this dataset is designed to capture authentic emotional expressions in real-world speech. As of Release 1.11, it comprises 151 654 utterances from more than 1500 speakers, totaling over 237 hours of speech.

Annotations are obtained via crowdsourcing, where each utterance is labeled by at least five annotators (with an average of 6.7 annotations per utterance). The labeling process includes both categorical emotion annotations (with a primary emotion and optional secondary emotions) of *anger*, *sadness*, *happiness*, *surprise*, *fear*, *disgust*, *contempt*, *neutral*, and dimensional ratings for *arousal*, *valence*, and *dominance* on a seven-point Likert scale. Majority voting is used to finalize categorical emotion labels, while dimensional ratings are determined by averaging annotators' scores.

To facilitate research and model evaluation, the dataset is partitioned into predefined training (84,030 utterances), development (19,815 utterances), test set 1 (30,647 utterances), and test set 2 (14,815 utterances) subsets. Given its diversity in speaking styles, topics, and recording conditions, MSP-Podcast serves as a crucial resource for developing SER models capable of handling natural, spontaneous speech.

CMU-MOSI

The *CMU Multimodal Opinion Sentiment Intensity* (CMU-MOSI) dataset (Zadeh et al., 2016) is a multimodal corpus designed for sentiment and emotion analysis in spoken language. It consists of monologues extracted from YouTube videos, where speech is accompanied by text transcriptions and video. The dataset comprises 2199 utterances from 93 speakers, covering a variety of topics and speaking styles.

Each utterance is annotated for sentiment polarity on a seven-point Likert scale (-3 = highly negative, +3 = highly positive). Due to its combination of audio, textual, and visual modalities, CMU-MOSI is widely used in multimodal emotion recognition research. However, since the speech samples originate from online videos, they often contain background noise and variations in recording quality, posing additional challenges for SER models.

CMU-MOSEI

An extension of CMU-MOSI, the *CMU Multimodal Opinion Sentiment and Emotion Intensity Extended* (CMU-MOSEI) dataset (Zadeh, Liang, Poria, Cambria, & Morency, 2018) significantly increases the dataset's scale and speaker diversity. It consists of 23,500 video segments from over 1,000 speakers, capturing a broad range of opinions and emotional expressions.

Similar to CMU-MOSI, CMU-MOSEI includes speech, transcriptions, and video, with each utterance annotated for sentiment on a continuous scale (-3 to +3) and six primary emotions (*happy, sad, angry, surprised, fearful, and disgusted*). Dimensional emotion ratings for *valence, arousal, and dominance* are also provided. Compared to CMU-MOSI, CMU-MOSEI offers increased linguistic and emotional diversity, making it a more comprehensive resource for training SER models in real-world settings. However, like its predecessor, it includes variations in background noise and recording quality due to its reliance on online videos.

Current studies on these datasets typically use speech foundation models with FFNNs or Transformers-based architectures on waveforms directly, providing state-of-the-art performance (Ando et al., 2021; Leem, Fulford, Onnela, Gard, & Busso, 2023; Pepino et al., 2021; Siriwardhana, Reis, Weerasekera, & Nanayakkara, 2020; M. Xu, Zhang, & Zhang, 2021).

2.5 Evaluation Metrics

Evaluation metrics are used to assess the performance of machine learning models and are selected according to the nature of the task: classification or regression. Below, we describe the commonly used metrics in each category, along with their mathematical definitions.

2.5.1 Classification Tasks

Accuracy

Accuracy measures the proportion of correctly predicted samples among all predictions.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (2.21)$$

where TP , TN , FP , and FN denote true positives, true negatives, false positives, and false negatives, respectively.

- **Unweighted Accuracy (UA):** Also known as the overall accuracy, this metric calculates the fraction of instances predicted correctly, without considering class distribution:

$$\text{Unweighted Accuracy} = \frac{\text{Total correct predictions}}{\text{Total number of samples}} \quad (2.22)$$

- **Weighted Accuracy (WA):** Weighted accuracy is computed by averaging the accuracy within each class, weighted by the number of samples in that class:

$$\text{Weighted Accuracy} = \sum_{c=1}^C \frac{n_c}{n} \cdot \frac{TP_c}{n_c} = \frac{1}{n} \sum_{c=1}^C TP_c \quad (2.23)$$

where C is the number of classes, n_c is the number of samples in class c , n is the total sample number, and TP_c is the number of correctly predicted samples in class c .

The distinction between weighted and unweighted accuracy is especially important in the presence of class imbalance. While unweighted accuracy gives the same weight to every instance regardless of class, weighted accuracy ensures that each class contributes proportionally to its size in the dataset.

Precision, Recall, and F1 Score

$$\text{Precision} = \frac{TP}{TP + FP} \quad (2.24)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (2.25)$$

$$\text{F1 Score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (2.26)$$

These metrics can be averaged across classes using either *macro* (unweighted average) or *weighted* (weighted by support) averaging to handle class imbalance.

2.5.2 Regression Tasks

Mean Squared Error (MSE)

MSE measures the average of the squared differences between the predicted and true values.

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (2.27)$$

where y_i and \hat{y}_i denote the ground truth and predicted value for the i -th sample, respectively.

Pearson Correlation Coefficient (PCC)

PCC measures the linear correlation between predicted and ground truth values.

$$\text{PCC} = \frac{\sum_{i=1}^n (y_i - \bar{y})(\hat{y}_i - \bar{\hat{y}})}{\sqrt{\sum_{i=1}^n (y_i - \bar{y})^2} \sqrt{\sum_{i=1}^n (\hat{y}_i - \bar{\hat{y}})^2}} \quad (2.28)$$

where \bar{y} and $\bar{\hat{y}}$ are the means of the ground truth and predicted values, respectively.

Concordance Correlation Coefficient (CCC)

CCC assesses the degree to which pairs of observations fall on the 45-degree line through the origin. It combines measures of both precision and accuracy.

$$\text{CCC} = \frac{2\rho\sigma_y\sigma_{\hat{y}}}{\sigma_y^2 + \sigma_{\hat{y}}^2 + (\mu_y - \mu_{\hat{y}})^2} \quad (2.29)$$

where ρ is the Pearson correlation coefficient, σ_y and $\sigma_{\hat{y}}$ are standard deviations, and μ_y , $\mu_{\hat{y}}$ are means of the ground truth and predicted values, respectively.

The choice of evaluation metrics depends on the nature of the learning task, the data distribution, and the annotation format of the labels. In this thesis, we primarily focus on **unweighted accuracy** for classification tasks, as it is the most commonly used metric in SER. For regression tasks, we mainly adopt **mean squared error** and **concordance correlation coefficient** when modeling continuous emotional dimensions. Several other metrics will occasionally be included when necessary.

2.6 Prerequisites for Integrating ASR into SER

As the prerequisite of integrating ASR into SER, the relationship between SER and ASR needs to be understood. This section describes the ASR task definition and ASR models used in this thesis, as well as introduces the background of SER-ASR relationship and the use of ASR for SER, leading to the thesis work.

2.6.1 ASR Task Definition

ASR aims to convert spoken language into written text by estimating the most probable word sequence \hat{W} given an acoustic input X . This is commonly formulated as a probabilistic inference problem:

$$\hat{W} = \arg \max_W P(W|X) \quad (2.30)$$

where W denotes a sequence of words and $P(W|X)$ is the posterior probability of the word sequence given the observed acoustic features.

However, directly modeling the posterior distribution is often challenging. A more tractable approach is to apply Bayes' theorem to decompose the posterior as follows:

$$P(W|X) = \frac{P(X|W)P(W)}{P(X)} \quad (2.31)$$

Since $P(X)$ is independent of W , it can be omitted during maximization. Therefore, the problem reduces to:

$$\hat{W} = \arg \max_W P(X|W)P(W) \quad (2.32)$$

In this Bayesian formulation:

- $P(X|W)$ is the *acoustic model*, which estimates the likelihood of the acoustic input given a word sequence.
- $P(W)$ is the *language model*, which provides a prior over plausible word sequences.

This decomposition enables separate modeling of acoustic and linguistic knowledge and forms the theoretical basis for conventional ASR systems.

In SER, ASR systems can provide two types of representations:

Acoustic representation: The internal hidden representations of ASR models (e.g., encoder outputs) can be used as high-level acoustic features that are often emotion-sensitive.

Lexical representation: The transcriptions generated by ASR systems provide textual input for text-based SER models, enabling semantic-level emotion inference.

2.6.2 Evaluation Metrics

The performance of ASR systems is commonly evaluated using error rates that measure the mismatch between the predicted transcription and the reference transcription. The most widely used metrics include:

Word Error Rate (WER) measures the proportion of words that are incorrectly predicted. It is calculated based on the minimum number of insertions (I), deletions (D), and substitutions (S) needed to convert the predicted sequence into the reference:

$$\text{WER} = \frac{S + D + I}{N} \quad (2.33)$$

where N is the total number of words in the reference.

Character Error Rate (CER) is similar to WER, but computed at the character level, which is particularly useful for languages with no explicit word boundaries or for low-resource scenarios:

$$\text{CER} = \frac{S + D + I}{N} \quad (2.34)$$

where N here refers to the number of characters in the reference.

To evaluate the quality of ASR transcriptions for downstream tasks such as SER, metrics from the machine translation domain that emphasize n -gram matching are also commonly applied:

BLEU (Bilingual Evaluation Understudy) is a metric used to evaluate the quality of machine-generated translations by comparing n -grams (contiguous sequences of n words) between the hypothesis (machine output) and reference translations (Papineni, Roukos, Ward, & Zhu, 2002). It primarily measures *precision* but incorporates a *brevity penalty* to avoid overly short translations:

$$\text{BLEU} = \text{BP} \cdot \exp \left(\sum_{n=1}^N w_n \cdot \log p_n \right) \quad (2.35)$$

Where:

- N is the highest n -gram level considered (typically $N = 4$).
- w_n is the weight for each n -gram level, usually set as $w_n = \frac{1}{N}$.
- p_n is the precision for the n -gram level, defined as the ratio of the number of matching n -grams between the hypothesis and reference to the total n -grams in the hypothesis.
- BP is the brevity penalty, given by:

$$\text{BP} = \begin{cases} 1, & \text{if } c > r \\ e^{(1-\frac{c}{r})}, & \text{if } c \leq r \end{cases}$$

where c is the length of the candidate translation and r is the length of the reference translation.

The BLEU score ranges from 0 to 1, with a score of 1 indicating perfect matching of n -grams between the hypothesis and reference(s). The brevity penalty helps to prevent systems from generating very short translations just to match a few words.

GLEU (Google BLEU) is a variant of BLEU that is designed for sentence-level evaluation, making it particularly useful for tasks such as text correction, grammar correction, and short-form translation (Y. Wu et al., 2016). Unlike BLEU, which primarily focuses on precision, GLEU takes both *precision* and *recall* into account:

$$\text{GLEU} = \min(\text{Precision}, \text{Recall}) \quad (2.36)$$

Where:

- *Precision* is the ratio of the number of matching n -grams in the hypothesis to the total number of n -grams in the hypothesis:

$$\text{Precision} = \frac{\text{Number of matching } n\text{-grams in hypothesis}}{\text{Total number of } n\text{-grams in hypothesis}}$$

- *Recall* is the ratio of the number of matching n -grams in the hypothesis to the total number of n -grams in the reference:

$$\text{Recall} = \frac{\text{Number of matching } n\text{-grams in hypothesis}}{\text{Total number of } n\text{-grams in reference}}$$

GLEU evaluates both precision and recall, ensuring that the hypothesis captures as many correct n -grams from the reference as possible while maintaining fluency. The final score is the minimum of the precision and recall values.

BLEU is more suited for evaluating larger corpora or multi-sentence outputs, as it primarily measures precision and applies a brevity penalty for shorter translations. GLEU, on the other hand, is optimized for sentence-level tasks and provides a more balanced evaluation by considering both precision and recall, making it ideal for tasks like grammatical error correction or speech recognition.

2.6.3 Relationship between SER and ASR

The relationship between ASR and SER is not yet clearly defined. Although both of them take speech as input, the acoustic model in ASR focuses on the phone and word level, while SER focuses on the utterance level and beyond. Previous studies showed that the presence of emotion in speech degrades the accuracy of ASR (Fernandez, 2004). On the other hand, prior work has reported improvement in SER accuracy when lexical input is added to the acoustic input (C. M. Lee & Narayanan, 2005).

Transferring ASR for SER

Computer systems do not have the complexity of the human brain and auditory system to perfectly handle multiple tasks simultaneously. However, most speech features can be shared for multiple speech tasks, such as MFCC and F0. Therefore, the same speech features can be used as input for different speech tasks where they may be then further tuned with distinct task characteristics by the model for specific tasks. For this reason, intermediate representations from a trained ASR model can be transferred for SER. These representations can then be used as-is or fine-tuned on given SER-labeled data.

In 2016, Fayek, Lech, and Cavedon (2016) demonstrated the possibility of knowledge transfer from a convolutional neural acoustic model trained on the TIMIT corpus (Garofolo, Lamel, Fisher, Fiscus, & Pallett, 1993) for the emotion recognition task. The authors provided many fine-tuning variants, gradually holding some layers constant and fine-tuning the remaining ones. When compared to an end-to-end CNN model, employing the ASR network as a feature extractor and training only the output softmax layer resulted in a considerable reduction in performance. Allowing more ASR layers to be updated during back-propagation resulted in gradual improvements, but overall, employing ASR for feature extraction had a negative impact on SER performance.

In 2018, Lakomkin, Weber, Magg, and Wermter (2018) compared two options of tuning ASR representations: simply feeding them into a softmax classifier or adding a new Gated Recurrent Units (GRU) layer trained on top of the ASR features. The latter allows the model to train a SER network parallel to the ASR network to capture additional emotion-specific information. The ASR branch of the network remains frozen in both options. In the same year, Tits, El Haddad, and Dutoit (2018) showed that features learned by a trained ASR network can be used for SER and outperform the eGeMAPS feature set (Eyben et al., 2015). They also examined the relationship of the emotional dimensions arousal and valence with the ASR layers of audio and text modalities. They showed that for some speakers, arousal is more correlated to hidden features extracted from the first three layers (closer to audio) and valence to the ones closer to the last three layers (closer to text). Then in 2020, Zhou and Beigi (2020) indicated potential feature overlap between ASR and SER, by showing that using representations of the 12th layer of a pre-trained Time-Delay Neural Network (TDNN)-based

ASR model for SER can achieve a performance comparable to state-of-the-art SER models. They also noticed a performance degradation in the layers closer to the final output layer since they are more specialized in speech recognition. Another work in the same year (Yeh, Lin, & Lee, 2020) first fine-tuned the Listen-Attend-Spell (LAS) model (Chan, Jaitly, Le, & Vinyals, 2016) on IEMOCAP and then adopted a Singular Value Decomposition (SVD)-based domain adaptation to remove redundant parameters of the pre-trained ASR model to adapt it to SER. Their method improved both the speech recognition rate and the emotion recognition accuracy on IEMOCAP. Furthermore, they also found representations from a deeper layer degrade the emotion recognition accuracy, suggesting that deeper layers of ASR may contain less emotional information.

There are also several work directly using ASR outputs as features for SER and compared them with ground-truth transcripts (Sahu, Mitra, Seneviratne, & Espy-Wilson, 2019a; Schuller, Batliner, Steidl, & Seppi, 2009), which also demonstrated the feasibility of incorporating ASR for SER. All these prior studies have provided solid experimental results and discussion, yet this direction still calls for more attention and detailed analysis, especially at the moment when speech foundation models are becoming mainstream, thus eliciting one of the topics of this thesis work: evaluating and utilizing speech foundation models for SER.

Joint Training ASR and SER

The other pathway to incorporate ASR in SER is multitask learning, also known as joint training. To note that, however, due to differences in corpora, data sizes, and modeling approaches, it is not currently feasible to train two tasks simultaneously from scratch. Therefore, joint training here means training the SER model and fine-tuning the pre-trained ASR model simultaneously. Past ASR systems hindered joint training with SER partly due to the need for additional modeling of the language model (as Formula 2.32). Thanks to the development of end-to-end ASR modeling in recent years (Amodei et al., 2016; Graves & Jaitly, 2014), joint ASR-SER training has become possible and as far as we know, it was firstly attempted in H. Feng, Ueno, and Kawahara (2020).

In this work, an end-to-end SER model combined with an acoustic-to-word ASR model (Soltau, Liao, & Sak, 2017) was proposed. The authors proposed an ASR feature, which is the hidden state of the decoder in the ASR model, to replace the textual output for the combination in the SER model. As a result, this is a speech-only ASR-SER model but has a performance that is close to that of the models using speech and text. By using multitask learning, the ASR model can be fine-tuned, and the SER model can get more accurate ASR features to improve the recognition performance. In 2021, Cai, Yuan, Zheng, Huang, and Church (2021) leveraged the pre-trained wav2vec 2.0 as the feature extractor backbone for joint ASR-SER training. They added a Connectionist Temporal Classification (CTC) head and a classification head on top

of the backbone, to simultaneously obtain predictions for transcripts and emotion categories. Moreover, they also discussed how ASR affects SER by adjusting the ASR weight in multitask loss. The evaluation results indicated that the ASR task is very important to SER, and a better trained ASR model (lower error rate) yields better SER performance.

2.7 Challenges in SER and ASR Integration

Despite the aforementioned research findings and significant advancements in feature engineering, model development, and dataset collection, SER continues to face several critical challenges that impede its progress. Addressing these challenges is crucial for developing more robust and generalizable SER systems.

2.7.1 Limited Availability of Emotional Speech Data

One of the most pressing challenges in SER research is the scarcity of large-scale emotional speech datasets. Unlike ASR, which benefits from extensive corpora such as LibriSpeech (Panayotov, Chen, Povey, & Khudanpur, 2015) ($\approx 1,000$ hours) and Common Voice (Ardila et al., 2020) ($\approx 33,535$ hours), the size of publicly available emotional speech datasets remains significantly smaller (e.g., IEMOCAP (Busso et al., 2008) ≈ 12 hours, MSP-Podcast (Lotfian & Busso, 2017) ≈ 237 hours). The collection of emotional speech data is inherently challenging due to the time-consuming and costly nature of the process. As a result, existing datasets often suffer from limited speaker diversity, constrained emotional variability, and a lack of varied speaking styles. This data limitation can lead to overfitting in SER models, reducing their generalizability, particularly in real-world scenarios where emotional expressions are highly diverse and context-dependent. To overcome this issue, there is a pressing need for large-scale, diverse, and publicly accessible emotional speech datasets that encompass a wide range of emotions, speaking styles, and recording environments.

2.7.2 Mismatch Between Speech Models and Emotional Speech

A consequence of data scarcity is the lack of speech models specifically designed for emotional speech processing.¹

Most conventional speech models, such as GMM-HMM models in Kaldi (Povey et al., 2011), Deep Speech (Hannun et al., 2014), and Jasper (J. Li et al., 2019), are optimized for ASR tasks, primarily focusing on phone and word recognition, rather than capturing emotional cues. As a result, adapting pre-trained ASR models for SER has proven to be challenging (Lakomkin, Weber, Magg, & Wermter, 2017; Yeh et al., 2020; Zhou & Beigi, 2020).

1. While a pre-trained model tailored for SER, `emotion2vec`, was introduced in 2023, it emerged after the commencement of this thesis project and is limited to SER tasks without ASR capabilities.

Recent speech foundation models, including wav2vec 2.0, HuBERT, and Whisper, have demonstrated impressive performance across various speech processing tasks, including SER. Unlike conventional handcrafted features or speech waveforms that generally work with sophisticated neural networks, the speech features produced from speech foundation models require only simple FFNNs to achieve state-of-the-art SER performance (Pepino et al., 2021; S.-w. Yang et al., 2021).

However, these models are not inherently designed to capture emotional characteristics in speech. First, their training are primarily conducted on non-emotional speech, lacking variations in prosody, voice quality, and speaking style. These aspects may not be fully captured by general-purpose speech models, potentially leading to distortions and suboptimal performance in SER. Second, the training targets of most self-supervised speech foundation models are divided into generative or discriminative approaches based on masked speech segmentation. These targets are mainly for reconstructing content from context (Chung, Hsu, Tang, & Glass, 2019; Chung, Tang, & Glass, 2020), pushing the model to generalize and potentially discard fine-grained paralinguistic cues.

2.7.3 Challenges with Lexical Information

Given the limitations of both conventional speech models and modern foundation models in handling emotional speech, another critical challenge arises: the reliance on imperfect lexical information.

In SER, lexical features complement acoustic cues to enhance emotion recognition. Traditionally, many SER systems have relied on human-annotated transcripts for lexical analysis. However, in real-world applications, such perfectly transcribed text is unavailable due to the absence of human annotators for immediate correction. Instead, ASR systems must be used to generate lexical information. Unfortunately, ASR models often struggle with emotional speech, as they are not explicitly trained to handle variations in prosody and expressiveness (Fernandez (2004)). For example, Goldwater, Jurafsky, and Manning (2010) found that words with extreme prosodic characteristics usually produce higher word error rates in speech recognition. These phenomena result in transcription errors that can negatively impact the accuracy of SER systems. Consequently, improving the robustness of SER in the presence of ASR errors remains a major research challenge.

2.7.4 Challenges with Integration of Audio and Text Modalities

Building on the previous challenge, the integration of audio and text modalities presents an additional hurdle for SER due to inherent discrepancies between these two modalities. While emotions in speech are conveyed through variations in acoustic properties, textual representations primarily rely on emotionally salient words. This fundamental difference makes it difficult to achieve seamless integration between the two sources of information.

Furthermore, ASR-generated transcriptions often contain errors, leading to misalignment between audio and text representations. Such inconsistencies can hinder multimodal fusion strategies and degrade SER performance. For example, Sahu, Mitra, Seneviratne, and Espy-Wilson (2019b) used two commercial ASR systems to generate transcripts for bimodal SER (audio + text), resulting in a relative loss of 4% and 5.3% in unweighted accuracy compared to ground-truth transcripts, respectively. Additionally, certain emotions may be conveyed exclusively through either audio or text, or may even exhibit contradictory expressions across modalities. These complexities further complicate the effective fusion of acoustic and lexical features, posing a significant challenge for multimodal SER research.

2.8 Chapter Summary

This chapter introduced the theoretical and practical foundations of SER and the background of integrating ASR into SER, covering task definitions, features, models, datasets, evaluation protocols, related literature, and key challenges. These insights lay the groundwork for the experiments and model proposals in subsequent chapters.

Utilizing Self-Supervised Acoustic Representations for SER

3.1 Introduction

The progression from traditional statistical ASR systems to advanced deep learning architectures and large-scale foundation models has greatly transformed the field of speech processing. In particular, these foundation models, while originally trained for ASR, have shown transferable potential across various speech-related downstream tasks, including SER. However, despite their success, the internal mechanisms by which these models capture, encode, and leverage emotional information are still not understood.

Unlike traditional speech modeling approaches that have been extensively researched, self-supervised learning (Self-SL) models have just started to be explored in very recent years, with wav2vec 2.0 (W2V2) attracting the most attention for its wide application potential. For example, Pasad, Chou, and Livescu (2021) conducted layer-wise analysis of W2V2 using a suite of tools and found 1) acoustic and linguistic properties are encoded in different layers; 2) the pre-trained model follows an autoencoder-style behavior; 3) the model encodes some non-trivial word meaning information. Fan, Li, Zhou, and Xu (2020) showed that W2V2 has the ability to discriminate between speakers and also languages, and this distinction is more obvious in lower layers. They hence proposed multi-task learning of speaker verification and language identification, and verified its feasibility. S.-w. Yang et al. (2021) set up benchmark performance using self-supervised speech models on a range of tasks.

Nevertheless, these Self-SL speech models are still understudied and the above-mentioned works have limitations. For example, Pasad et al. (2021) did not extend their exploration to emotional speech and SER tasks. In Fan et al. (2020), only a portion of the layer difference was shown, so misses a thorough layer-wise analysis. In S.-w. Yang et al. (2021), they presented downstream task performance without further explanation. Furthermore, none of those studies investigated paralinguistic characteristics in W2V2 representations. As such, in the thesis work, we build on previous work while adding new perspectives from detailed quantitative analysis on emotional corpora.

In this chapter, we aim to shed light on how and what speech foundation models contribute to SER, examining their behavior from the acoustic perspective. We perform a series of in-depth experiments to analyze the types of emotional cues embedded in the representations learned by these models. Specifically, we investigate the paralinguistic characteristics, hierarchical structures, and emotion biases, with the goal of understanding how to probe foundation models that are pre-trained upstream for downstream speech tasks like SER. Furthermore, we propose a hierarchical joint ASR-SER training model, integrating W2V2 outputs in SER, aiming to increase the SER accuracy and the reliability of ASR models on emotional speech.

The remainder of this chapter is structured as follows. Section 3.2 details the probing experiments and results examining W2V2 from an acoustic perspective on emotional speech. Section 3.3 describes a hierarchical attention-based fusion model that incorporates W2V2 outputs into joint ASR-SER training. Finally, Section 3.4 provides a summary of the chapter.

3.2 Exploring wav2vec 2.0 on Emotional Speech

Current speech foundation models often operate as black boxes when processing emotional speech, leaving the mechanisms by which they capture and represent emotional cues largely unexplored. Since there are various speech foundation models in terms of training schemes, number of layers, training data size, and more, in this section, we focus on W2V2 series, considering that W2V2 is the most popular speech foundation model and the less well-known black-box characteristics of Self-SL models, to frame our probing experiments and pave the way for probing other speech foundation models.

We begin by evaluating the SER performance of W2V2, aiming to reveal the extent to which it learned representations contribute to effective emotion classification. Following this, we conduct a detailed analysis of their paralinguistic abilities to determine whether it preserves key paralinguistic features known to correlate with human emotional perception, such as pitch, energy, and spectral characteristics.

In addition, we explore the hierarchical nature of W2V2 by probing different layers to identify which levels of their learned representations are most informative for SER tasks. This layer-wise analysis sheds light on whether emotional cues are captured at early (closer to the input), mid-level, or deeper (closer to the output) layers.

Finally, we examine the presence of emotion bias in W2V2, investigating whether their internal representations treat emotional content in a neutral, emotion-agnostic manner, or whether they inherently encode emotion-aware features.

Together, these analyses offer valuable insights into the interpretability of speech foundation models and their potential for emotion recognition applications.

3.2.1 Experimental Setup

Datasets

The IEMOCAP (IEM) and RAVDESS (RAV) datasets are used in this section. For IEM, following prior research (Y. Li, Zhao, et al., 2019; Peng, Zhu, Unoki, Dang, & Akagi, 2018), we combined *Happy* and *Excited*, and removed utterances that do not have transcriptions, bringing the total number of utterances used in this study to 5 500, each with one label from four classes: *Angry*, *Happy*, *Neutral*, and *Sad*. For RAVDESS, we only use the speech set, which has 1 440 utterances from 24 actors (12 female, 12 male) in eight emotions: *Calm*, *Happy*, *Sad*, *Angry*, *Fearful*, *Surprised*, *Disgust*, and *Neutral*. The major reason that we choose to use RAVDESS is that, even though other corpora may have a larger size, it provides fixed sentences with different emotional expressions. Such a setting excludes the lexical influence by “forcing” different emotions to have the same linguistic content, thus helping us to better explore the effects of the acoustic properties of W2V2 by eliminating the effects raised by lexical content (e.g., word pronunciation causing prosody variation).

Speech Foundation Models

We use *wav2vec2-base*, *wav2vec2-base-100h*, and *wav2vec2-base-960h* models, which are the pre-trained and fine-tuned models (on 100h and 960h of Librispeech) respectively. We refer to them as *PT*, *FT100*, and *FT960*. As mentioned before, we choose W2V2 because it is the most widely used speech foundation model (particularly one of the black-box Self-SL models), with the expectation that the exploratory approach can be generalized to other models.

3.2.2 Probing SER Performance

We first implement a layer-wise analysis by using the output of every individual layer within the Transformer network to demonstrate how information encoded by W2V2 contributes to SER. Next, as there is no common practice of how to utilize W2V2 representations as input features for downstream tasks, we compare the performance of three commonly used approaches of using W2V2 representations as input features:

- 1) taking the last layer output (Cai et al., 2021; L.-W. Chen & Rudnicky, 2023; Sharma, 2022);
- 2) taking the average of all layer outputs (Boigne, Liyanage, & Östrem, 2020);
- 3) taking the weighted average of all layer outputs (assigning a trainable weight to each layer output) (Pepino et al., 2021; S.-w. Yang et al., 2021).

We also propose a fourth approach which excludes the last two layers from averaging as they generally underperform other layers. We evaluate the performance using Unweighted Accuracy (UA). Like most downstream tasks, we use W2V2 models as frozen feature extractors.

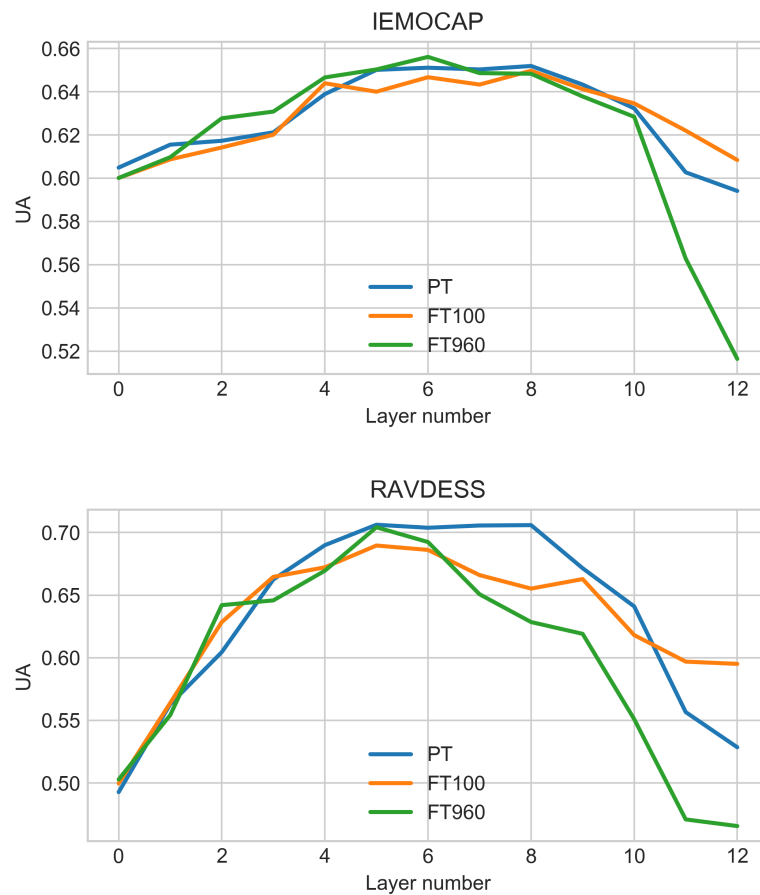


Figure 3.1: SER accuracy comparison using models. PT: pre-trained; FT100: fine-tuned on 100h of Librispeech; FT960: fine-tuned on 960h of Librispeech.

Since our goal is to explore information in W2V2 representations, we build a simple downstream model comprising only two dense layers (128 and 16 neurons, respectively) with *ReLU* activation and one output layer (four neurons for IEM and eight neurons for RAV) with *Softmax* activation. The learning rate is set as 1.0×10^{-4} and 2.0×10^{-4} for IEM and RAV with the *AdamW* optimizer, respectively, and the weight decay is set as 1.0×10^{-5} . The batch size is 64, and we train the models until validation loss converges, as different layer outputs converge at different steps. For IEM, we implement 5-fold Cross-Validation (CV) in accordance with prior work (Y. Li, Zhao, et al., 2019). For RAV, we randomly divide 24 speakers into four groups and implement 4-fold cross validation.

Figure 3.1 depicts the trends of layer-wise UA on the two corpora. We include layer 0 (the output of the CNN encoder right before the Transformers) in accordance with prior works. We see that:

Table 3.1: UA (%) using different inputs and models.

Input	Model	Corpus	
		IEM	RAV
Best layer	<i>PT</i>	65.19	70.62
	<i>FT100</i>	64.97	68.96
	<i>FT960</i>	65.61	70.42
Last layer	<i>PT</i>	59.41	52.85
	<i>FT100</i>	60.84	59.51
	<i>FT960</i>	51.64	46.56
Average	<i>PT</i>	64.93	67.26
	<i>FT100</i>	64.72	67.40
	<i>FT960</i>	65.51	64.20
Average w/o last two	<i>PT</i>	65.11	67.36
	<i>FT100</i>	64.90	67.50
	<i>FT960</i>	65.87	65.56
Weighted average	<i>PT</i>	65.28	68.47
	<i>FT100</i>	64.94	68.89
	<i>FT960</i>	65.67	65.11

1. Before the best middle layer (layer 6 for IEM and layer 5 for RAV), all three models (*PT*, *FT100*, and *FT960*) show the same trend: accuracies go up and are relatively close, but then start dropping after the middle layer. This upward-downward trend is possibly related to the acoustic-linguistic property of W2V2: frame-level inputs are encoded by the Transformers until the middle layer (Pasad et al., 2021). At this point, the representations encode phonetic information but have not yet lost much of the original acoustic properties. This makes the middle layers contain the most useful information for SER. In subsequent layers, the representations gradually encode word identity and word meaning more strongly (Pasad et al., 2021). At this stage, potential ASR errors together with the loss of the original acoustic information appear to lead to drops in SER accuracy.
2. On IEM, there are barely any differences among the UAs until layer 11, while on RAV, the differences after the middle layer are more dramatic. This phenomenon is plausible as RAV only has two fixed statements, yet IEM contains various sentences, allowing fine-tuned W2V2 models to make more use of linguistic information, which makes up for the acoustic loss. In RAV, however, every sentence is repeated with each emotion, which means linguistic information has no contribution to emotion discrimination.
3. In general, $PT > FT100 > FT960$ from the middle layer but *FT100* clearly outperforms the other two on the last two layers. We assume that moderate fine-tuning enables W2V2 to achieve a good acoustic-linguistic balance, as word information has been found encoded by the last two layers in fine-tuned models (Pasad et al., 2021), and such a balance helps *FT100* achieve better performance on the last two layers.

Table 3.1 compares the SER accuracies using different inputs and models, and yields the following findings:

1. UAs of the best layer of *PT* and *FT960* are close and higher than *FT100*. So, although moderate fine-tuning enables the model to capture both acoustic and linguistic properties for *FT100*, neither of them is fully encoded causing a decrease in accuracy for the best layer of *FT100*, compared to *PT* and *FT960*.
2. The situation is reversed on the last layer. Compared to *FT100*, *PT* lacks linguistic information and *FT960* relies too much on imperfect linguistic information while losing acoustic information due to “over” fine-tuning.
3. Word-level information does not help SER on RAV, as mentioned before, which makes the deeper layers of *FT960* the worst. Hence, it is reasonable that *FT960* generates better performance on IEM yet worse performance on RAV when taking the average on all layers, the average without the last two layers, or the weighted average on all layers as input.
4. In corpora like RAV, where emotional cues are conveyed only through speech acoustics, the best middle layer may suffice as input for SER. However, for corpora like IEM, emotional expression is also influenced by linguistic content.
5. Except for the “best” layer inputs, *FT960* and *FT100* produce the best results on IEM and RAV, respectively. This differs from the patterns in Figure 3.1 from which we would expect *FT960* to perform the worst on IEM and *PT* the best on RAV. It means that the performance obtained by averaging layer outputs does not equal the average of all layer performance, which demonstrates that representations of different layer contain different information contributing to SER.

3.2.3 Probing Paralinguistic Properties

Compared to handcrafted features, Self-SL representations are hard to explain as they are not designed by human knowledge of speech signals. For example, paralinguistic features have proven useful and have long been used for emotion detection or clinical speech analysis (Haider, De La Fuente, & Luz, 2019). On the other hand, Self-SL representations are formed by reconstruction loss or contrastive loss from speech, but whether, why, and how these losses contribute to downstream tasks other than ASR are still unknown.

In this experiment, we measure the similarities between each layer’s output and different types of paralinguistic features to see how W2V2 retains well-known acoustic correlates of speech perception. We evaluate the similarity using Canonical Correlation Analysis (CCA) (Hardoon, Szedmak, & Shawe-Taylor, 2004). CCA finds linear projections of two multivariate feature sets such that the correlations between the projected variables are maximized. It is particularly

suitable for our task because it can capture shared information between the model representations and paralinguistic features even if they exist in different representational spaces. By using CCA, we can quantitatively assess how closely the internal representations of W2V2 align with known acoustic correlates.

The selected paralinguistic features are listed in Table 3.2, which are mainly based on eGeMAPS (Eyben et al., 2015), commonly used as a minimal set of features for SER. The groupings are directly adopted from eGeMAPS, except that Shimmer and Jitter are grouped under voice quality, as they have proven especially useful yet are often overlooked in SER (Jacob, 2016; X. Li et al., 2007). We also extract 40 MFCCs as linguistic (phone) features for comparison. We downsample them to make their sequence lengths comparable to W2V2 representations as required by CCA.

Figure 3.2 shows the layer-wise CCA on IEM and RAV using three W2V2 models. It can be noted that:

1. The curves of paralinguistic features are flatter than those of MFCCs, indicating that W2V2 does not focus on as much paralinguistic information. In particular, voice quality features seem barely taken into the encoding process.
2. When looking at *PT* models, we can see a reverse trend from the middle layer, which reinforces a view that the *PT* model follows an autoencoder style behavior where deeper layers “reconstruct” the input (Pasad et al., 2021). The reverse trend of similarity with MFCCs on IEM is weaker than that on RAV, possibly demonstrating that the linguistic complexity makes the reconstruction process harder and more error-prone. Hence, the deeper representations are more similar to MFCCs on RAV than on IEM. The peculiar pattern of the last two layers is due to the training objective of masked segment prediction.
3. When looking at the fine-tuned models, we note that the similarities keep decreasing since the models have been fine-tuned towards ASR and learn to compute speech information from frame to phone, and then to word level with layer depth (Pasad et al., 2021). This phenomenon reinforces our explanation of the accuracy drop in Figure 3.1 that acoustic properties are being replaced by linguistic ones that contain errors.
4. The graphs indicate that the overall similarity variation on IEM is larger than on RAV. This is again, likely due to the fact that RAV has much less linguistic variation overall, which we in turn see as less change in CCA through the layers. Moreover, since how we say a sentence is affected by its content, the CCA variation of paralinguistic features is larger on IEM than on RAV. Finally, since the layer outputs contain more complex linguistic information, the overall CCA values for paralinguistic features on IEM are

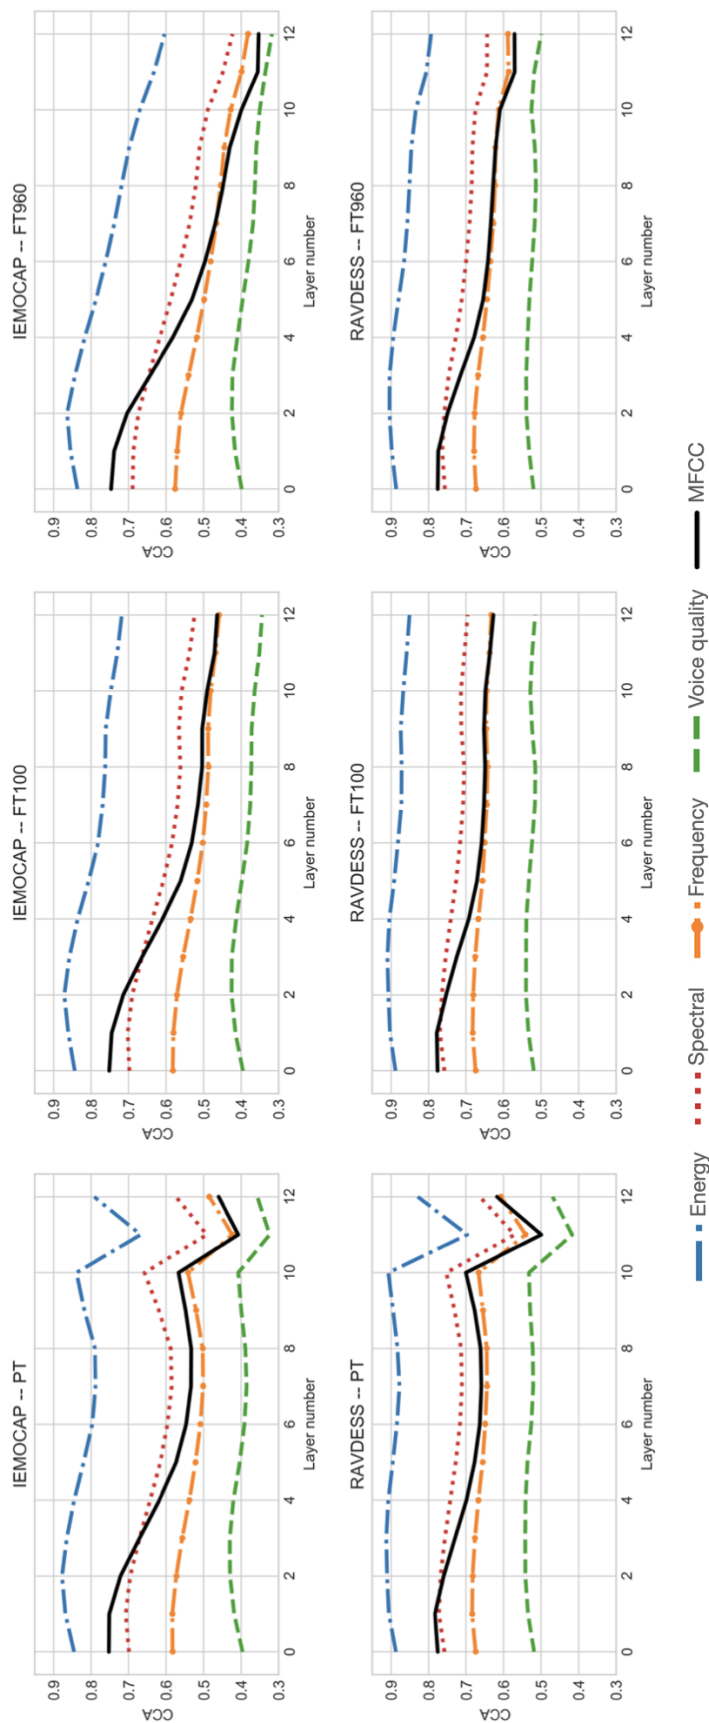
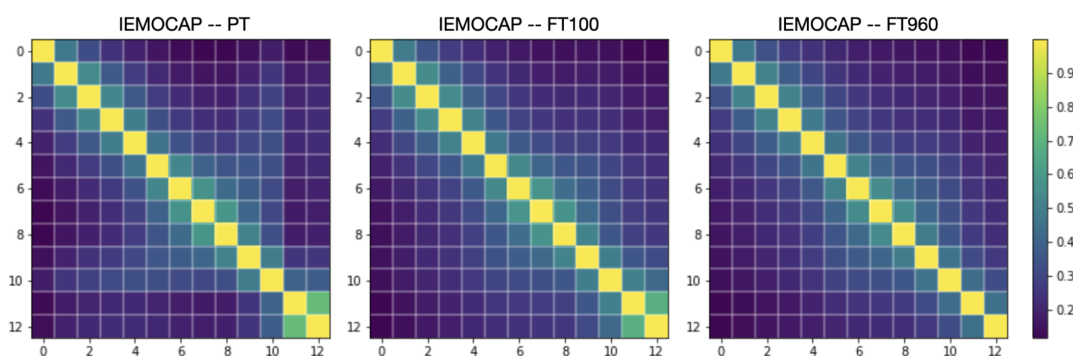


Figure 3.2: CCA similarity comparison for paralinguistic property.

Table 3.2: Extracted paralinguistic features.

Feature set	Low-level descriptors
Energy	Loudness; Harmonics-to-noise ratio
Frequency	Pitch; Formant 1; Formant 1, 2, 3 frequency
Spectral	Alpha ratio; Hammarberg index; Formant 1, 2, and 3 relative energy; Spectral slope 0-500 Hz, 500-1500 Hz; Harmonic difference H1-H2 and H1-A3
Voice quality	Jitter; Shimmer

**Figure 3.3:** Pair-wise correlations of layer representations.

lower than those on RAV, no matter the starting values or overall values. However, the starting values of similarities with MFCCs on both corpora are almost the same, further suggesting that W2V2 focuses on learning linguistic information rather than paralinguistic.

3.2.4 Probing Layer Correlations

To better understand how different layer outputs are correlated with each other before and after fine-tuning for ASR, and how W2V2 encodes information and contributes to SER, we calculate pair-wise CCA similarities of W2V2 representations from every layer and plot the similarities using heat maps to visualize the correlations. We only discuss IEM, as the same patterns are found on RAV.

Figure 3.3 shows the pair-wise correlations:

1. In the *PT* model, we can notice an arrow-like shape from layers 0–10. Specifically, layers 9 and 10 show higher correlations with shallow layers compared to those in *FT100* and *FT960*. This suggests that representations in these layers become more general, making certain shallow-layer information easier to extract, which also explains the reverse trend observed in the *PT* model. However, such a pattern seems specific

Table 3.3: Correlations between the last two and prior layers.

Layer	Model	Layer 6	Layer 7	Layer 8	Layer 9	Layer 10
11	<i>PT</i>	0.25	0.24	0.24	0.31	0.45
	<i>FT100</i>	0.29	0.30	0.31	0.38	0.53
	<i>FT960</i>	0.24	0.26	0.27	0.32	0.43
12	<i>PT</i>	0.26	0.25	0.26	0.31	0.43
	<i>FT100</i>	0.27	0.28	0.29	0.34	0.44
	<i>FT960</i>	0.22	0.22	0.23	0.26	0.32

to these two layers as it is not obvious in prior deep layers, even layer 8, which also accounts for the reverse trend in Figure 3.2. After fine-tuning, this phenomenon disappears. The two layers become the same as shallow layers that have high correlations with nearby layers, and the correlations get weaker with distance.

2. In *PT* model, the last two layers are highly correlated with each other, even more obvious than the prior adjacent layers. However, the correlation gets weaker as fine-tuning goes. The color becomes dimmer in *FT100* and further dimmer in *FT960*. Nevertheless, an interesting phenomenon appears in *FT100*: the similarities between last two layers (especially layer 11) and the prior deep layers (layer 6 to 10) become higher. We present the CCA values in Table 3.3. It is obvious that the correlations in *FT100* are the highest, but they decrease in *FT960*. Moreover, the lower right part is slightly brighter in *FT100* than in *FT960* (values are skipped as limited space) indicating the deeper layers are more correlated with each other, which means low-level linguistic information (e.g., phonetics) generally exists. These phenomena validate our assumption that moderate fine-tuning enables W2V2 to achieve a good acoustic-linguistic balance but over fine-tuning “forces” the model to concentrate on learning high-level linguistic properties (e.g., word meaning) towards ASR.

3.2.5 Probing Hierarchical Properties

Speech generally follows the hierarchy of frame, phone, syllable, word and utterance from low to high level, that contribute differently to speech perception and understanding (Pascual, Ravanelli, Serra, Bonafonte, & Bengio, 2019; Yenigalla et al., 2018). Since Self-SL enables frames to capture context information, the representations are expected to contain higher-level meanings. To verify this, we prepare the extracted paralinguistic features at frame, phone, and word levels and measure their similarities with W2V2 representations using CCA, respectively. As our purpose is only to verify whether W2V2 features contain high-level speech information, we do not use forced alignment to determine the perfect boundaries. Instead, we adopt a less accurate yet efficient approach to compute hierarchical features: we approximate a phone by averaging five consecutive frames (we also tried to add overlap, but the results didn’t make

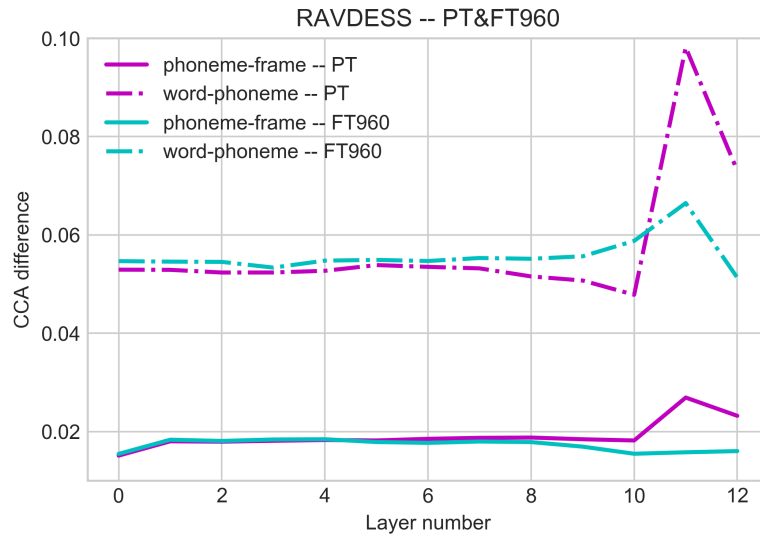


Figure 3.4: Hierarchical CCA similarity differences.

much difference): a word by averaging five consecutive phones, based on the fact that frame length is set as 25 ms when being extracted, and phone length varies from 50 ms to 200 ms (five frames on average) and word length from 250 ms to 1000 ms (five phones on average) (W. Chen, Xing, Xu, Pang, & Du, 2022).

We use all the paralinguistic features provided by eGeMAPS and implement the composition of hierarchical features. Then we downsample the paralinguistic features or the W2V2 representations depending on their lengths to make them comparable. Finally, we compute the CCA differences ($CCA_{phone} - CCA_{frame}$ and $CCA_{word} - CCA_{phone}$) which represent how similar the higher-level features with W2V2 representations are compared to the lower-level ones.

From Fig 3.4, we can note that:

1. Higher-level paralinguistic features do have higher similarities with W2V2 representations as the difference values are all positive. Besides, the value of $CCA_{word} - CCA_{phone}$ is even higher than $CCA_{phone} - CCA_{frame}$, which means W2V2 representations are more similar to word-level paralinguistic information.
2. The CCA differences barely change until layer 11 and become larger in the last two layers, which is due to the masked segment prediction enabling them to capture more context information (which is high level), especially on layer 11.
3. The curves of fine-tuned models are flatter because the last two layers become more coherent with the other layers by fine-tuning. Note that since the paralinguistic property is affected by linguistic property in IEM, the patterns are not as clear as RAV. Therefore, yet we observed similar trend in IEM, we only illustrate it using RAV.

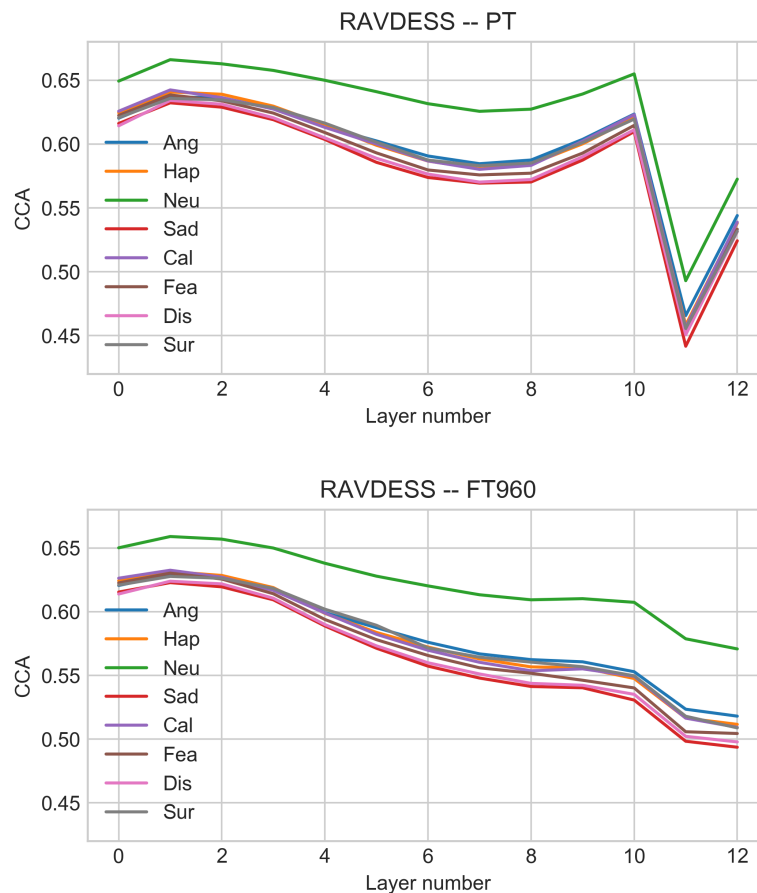


Figure 3.5: Discriminative analysis for emotion bias.

3.2.6 Probing Emotion Bias

Different emotions have different paralinguistic patterns (X. Li et al., 2007; Lugger & Yang, 2007). For example, hot angry and happy emotions usually have high intensity and pitch, while sad and calm emotions have low intensity (Luengo, Navas, & Hernandez, 2010). Hence, we calculate CCA similarities between paralinguistic features with W2V2 representations of every emotion for discriminative analysis. We also use all the paralinguistic features in eGeMAPS as in the previous task.

As illustrated in Fig 3.5:

1. Higher similarities between paralinguistic features with W2V2 representations are found in *Neutral* emotion for both the *PT* model and the *FT* models, pointing to interesting observations: a) *Neutral* is likely more frequently represented within Librispeech, as it is a corpus of read audiobooks where most emotional cues arise only within speech of fictional characters, i.e. bias in the data during pre-training of W2V2 consequently results in learned representations which are emotion-agnostic; b) the pre-training pretext

task in W2V2 (predicting masked segments) is not sufficient to learn a truly generalized representation in which different emotions are captured effectively. We also see that the curves converge after the middle layer on *PT* model. This again indicates that the deeper layers (except the last two) of *PT* model reconstruct the acoustic input.

2. The curves become even less distinguishable at the last two layers, indicating again the autoencoder type of learning resulting from the masked segment prediction does not help distinguish emotions. This may be because the paralinguistic information of a masked segment is difficult to predict from unmasked segments, as paralinguistic information is more spontaneous and less contextual compared to linguistic information. The masked segment prediction discriminates linguistic information while blurring the paralinguistic difference among frame segments, which makes the paralinguistic properties of every emotion become similar, resulting in the close curves.
3. For the *FT960* model, the distances between the curves increase with depth. It seems that W2V2 not only avoids encoding paralinguistic information, but consistently discards some paralinguistic features as learning proceeds, especially in speech that contains rich paralinguistic information, e.g., voice quality, which leads to increasing differences in the similarities between W2V2 representations and paralinguistic features across emotions (otherwise, the distances should not change with layer depth).

3.2.7 Summary

As speech foundation models are usually used simply as feature extractors for SER, there is a lack of understanding and approaches to analyze these models on emotional speech. To this end, we study W2V2, a representative speech foundation model, demonstrating how state-of-the-art ASR models should be explored for SER. We found that representations from W2V2 created for ASR lacks certain paralinguistic information. We also contribute to understanding the types of representations W2V2 learns by thoroughly comparing layer outputs in their correlations and SER. The hierarchy and bias analysis pave the path for better usage of W2V2 on downstream tasks. Our endeavor contributes to bridging the long-existing gap between ASR and SER research.

3.3 Hierarchical Fusion with Joint ASR Fine-Tuning

Considering the satisfactory SER accuracy achieved by W2V2 and its limited emphasis on paralinguistic features, as shown in the previous section, we propose using a single W2V2 model as the backbone, jointly fine-tuning it on emotional speech while simultaneously training a SER model.

Although both ASR and SER use speech signals as input, ASR operates more at the frame level, whereas SER recognizes emotion on larger timescales. Previous work (Fayek et al., 2016) has demonstrated that hidden features from the initial layers of both ASR and SER tasks are transferable, with their relevance gradually fading through deeper layers. Yoon et al. (2018) used the Google Cloud Speech API to generate transcription that were then fused with MFCCs for SER. They achieved 69.1% WA, likely due to their low WER (5.53%). (Sahu et al., 2019b) concatenated GloVe features obtained from Wit.ai API text transcription with acoustic features, obtaining a 62.9% UA using an LSTM based model. H. Feng et al. (2020) jointly trained ASR-SER using log mel-scale filter bank outputs and decoder outputs from a pre-trained ASR model, and achieved 68.6% WA. Cai et al. (2021) used a single W2V2 as the training model for both ASR and SER rather than separating two tasks, and used 10-fold cross-validation, which usually yields better results than 5-fold; Zhou and Beigi (2020) fine-tuned the pre-trained ASR model on IEMOCAP before transferring the ASR features for SER.

However, it remains unclear which ASR features benefit SER and in what way. To address this, we compare layer-wise ASR hidden features and ASR text outputs, and propose a hierarchical fusion of both into the pipeline for joint SER training.

3.3.1 Experimental Setup

For the **dataset**, we use IEMOCAP and removed utterances that did not have transcription, bringing the total number of utterances used in this study to 5,500. For the **ASR model**, we use the “wav2vec2-base-960h” model. Performance is evaluated using Unweighted Accuracy (UA).

For our proposed model, the audio time-series is encoded by the W2V2 model for ASR, and separately as MFCCs or the paralinguistic feature set from eGeMAPS (Table 3.2) for SER. On the ASR path, the W2V2 representations are then decoded to text by a word-level Connectionist Temporal Classification (CTC) decoder. Both the hidden state output and the text are extracted. A RoBERTa model (Y. Liu, 2019) is then used to extract lexical features from text output. On the SER path, a max pooling with kernel size 2 is conducted on the acoustic features (MFCCs or eGeMAPS) to obtain an acoustic representation.

The ASR hidden output (short for hidden layer output), text output, and pooled acoustic features, are then encoded using 2-layer Bidirectional Long Short-Term Memory (Bi-LSTM) networks. Each layer contains 32 hidden units followed by a dropout with probability of 0.5. A self-attention layer with 16 heads and 64 nodes is applied on the output of the Bi-LSTM, which generates a fixed-length vector as the feature encoding.

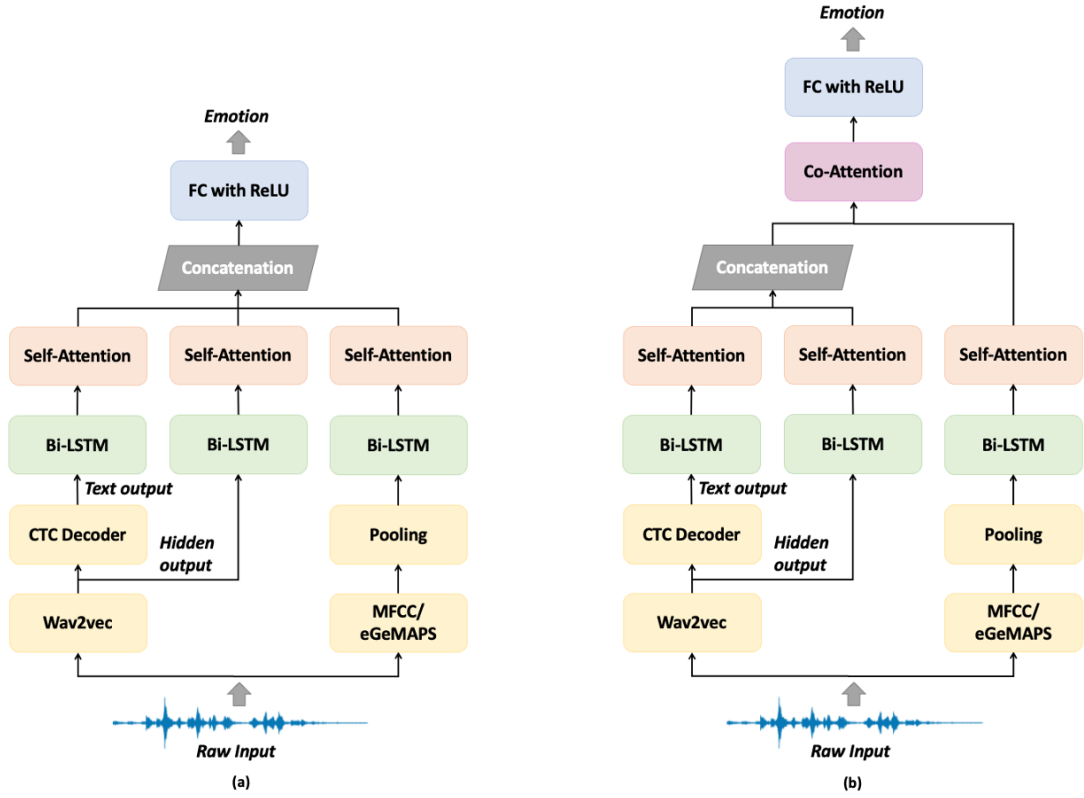


Figure 3.6: Baseline models. (a): concatenation; (b): concatenation with co-attention fusion.

3.3.2 Feature Fusion and Emotion Output

We fuse the three encodings from the self-attention layer to produce a final vector for emotion classification. We compared our model with three baseline models: **(a) concatenation** and **(b) concatenation with co-attention fusion**, as shown in **Figure 3.6**. Our proposed **hierarchical co-attention fusion** is shown in **Figure 3.7**. The co-attention mechanism (J. Lu, Batra, Parikh, & Lee, 2019) concatenates two hidden-state vectors which exchange key-value pairs in self-attention, allowing features from one input channel to be incorporated into the other:

$$H_C = \text{Concat}(H_{C_1}, H_{C_2}) \quad (3.1)$$

$$H_{C_i} = \text{MultiHead}(Q_A, K_B, V_B)W^O \quad (3.2)$$

$$= \text{Concat}(\text{head}_1, \dots, \text{head}_n) \quad (3.3)$$

$$\text{head}_i = \text{Attention}(Q_A W_i^Q, K_B W_i^K, V_B W_i^V) \quad (3.4)$$

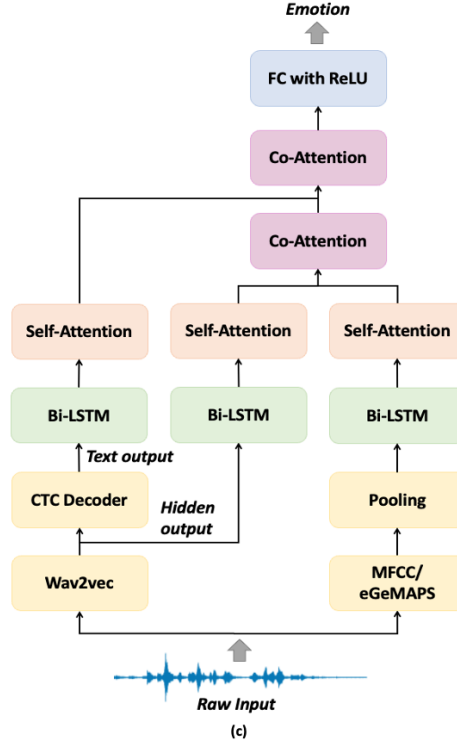


Figure 3.7: The proposed model. (c): hierarchical co-attention fusion with MFCC or eGeMAPS paralinguistic features as the acoustic input.

where W^O , W_i^Q , W_i^K , and W_i^V are trainable parameters. Q_A represents the query from one input channel, while K_B and V_B represent the key and value from the other. The value of n is 16, and H_C denotes the final concatenated hidden states of co-attention.

Hierarchical approaches for emotion recognition have proven useful in fusing different-level or different-type of features in previous works (Tian, Moore, & Lai, 2016). Inspired by these works and the hierarchical characteristics of speech (Pascual et al., 2019), we propose to fuse the input features from low to high level in a hierarchical manner using co-attention. Because MFCC and eGeMAPS paralinguistic features are extracted within frames, W2V2 features contain within-utterance context, and RoBERTa features can learn cross-sentence context, we refer to them as frame-level, utterance-level, and corpus-level respectively, and hierarchically fuse them in this order to generate a fixed-length vector. This is passed to a fully connected output layer with *Softmax* activation function to generate the probability distributions over emotion classes.

The model is optimized by the multi-task loss function:

$$\mathcal{L} = \lambda \mathcal{L}_{\text{ASR}} + (1 - \lambda) \mathcal{L}_{\text{SER}}, \quad (3.5)$$

Table 3.4: Performance comparison of using different models, features and fusion methods. $\lambda = 0.1$.

Model	Feature	Fusion approach	UA
1) SER (<i>baseline</i>)	MFCC + Ground-truth transcription	Concatenation	66.5%
		Co-attention	67.3%
	eGeMAPS + Ground-truth transcription	Concatenation	67.4%
		Co-attention	68.4%
2) ASR-SER	MFCC + Middle layer reps.	Concatenation	63.9%
		Co-attention	64.4%
	eGeMAPS + Middle layer reps.	Concatenation	66.9%
		Co-attention	67.7%
3) ASR-SER	MFCC + ASR text output	Concatenation	63.7%
		Co-attention	64.0%
	eGeMAPS + ASR text output	Concatenation	66.7%
		Co-attention	67.0%
4) ASR-SER	MFCC + Middle layer reps. + ASR text output	Concatenation	64.1%
		Co-attention	64.5%
		Hierarchical co-attention	66.1%
	eGeMAPS + Middle layer reps. + ASR text output	Concatenation	67.1%
		Co-attention	67.5%
		Hierarchical co-attention	68.4%

where \mathcal{L}_{ASR} and \mathcal{L}_{SER} are the losses for ASR and SER, respectively. We vary λ from 0.1 to 0.9 in increments of 0.2 to examine whether ASR or SER should serve as the primary task in the joint training framework. We use cross-entropy as the loss function and Adam optimizer with a learning rate of 1.0×10^{-4} and decay rate of 1.0×10^{-4} . The gradients are clipped with a threshold of 5.0. The batch size is set to 20, and the number of epochs is limited to 100. If the training of either task (ASR or SER) converges (i.e., its training loss does not decrease for three consecutive epochs), we terminate its training, allowing the joint model to focus on the other task. We perform 5-fold CV and use UA to assess the performance¹.

3.3.3 Results

The experimental results are presented in Table 3.4. We set $\lambda = 0.1$, as it provides the best performance. The results include: **1)** the baseline model, which combines acoustic features with human transcription in the SER model (i.e., without ASR); the joint ASR-SER model comprising **2)** acoustic features + ASR hidden output; **3)** acoustic features + text output; and **4)** our proposed full model. We utilized the middle-layer ASR output from the W2V2 model, as it achieves the best SER performance according to Table 3.1.

We notice that:

1. Note that the experimental settings have been revised from those in our published work (Y. Li, Bell, & Lai, 2022) and resulted in improved results.

1. Both ASR hidden output and text output help improve the SER performance over the acoustic features, although the difference is small with ASR hidden output slightly better than ASR text output. This is consistent with a previous claim in the literature that “ASR features can be more robust than the text output of ASR” (H. Feng et al., 2020).
2. In general, the performance of co-attention fusion outperforms that of concatenation. It appears that the relatedness of two input channels is learned by attention.
3. The eGeMAPS paralinguistic features contributes slightly more than MFCCs when combining with ground-truth transcription. However, when combining with the middle layer representations, the difference becomes larger. This is plausible as W2V2 representations lack some paralinguistic information as shown in Figure 3.2 in the previous section. Thus, combining both makes the features more comprehensive for SER.
4. Our proposed approach: the joint ASR-SER model incorporating both ASR hidden output and text output using hierarchical co-attention achieves the same UA as the eGeMAPS + ground-truth co-attention result. This indicates that our approach can help ameliorate ASR errors by combining paralinguistic features, hidden ASR features, and ASR text output.

To better understand the internal representations of the W2V2 model and their relevance to SER, we conducted a detailed analysis focusing on the contribution of different W2V2 layers. While Chapter 3.2 presented a layer-wise study on SER using W2V2 features, here we take a look into how it functions with joint ASR fine-tuning by comparing the hidden-state outputs from each layer. As shown in Table 3.5, the results align with the findings in Figure 3.1 where the SER accuracy increases till the middle layer then drops.

Additionally, the overall SER performance in this joint training setup is noticeably lower than the results reported in Figure 3.1, where SER was performed independently without simultaneous ASR tuning. We hypothesize that this degradation stems from the impact of the ASR training in the joint optimization process. In other words, since the model architecture is built upon W2V2, the emotional characteristics (such as voice quality features) in the representation space may be suppressed because it is primarily fine-tuned for speech recognition. This phenomenon highlights an inherent trade-off in multitask learning setups, where optimizing for one task, in this case, word recognition, may inadvertently diminish the model’s capacity to capture fine-grained emotional cues.

To verify our hypothesis, we investigate the impact of the weight of the ASR training λ using the hierarchical co-attention model proposed, as it performs best. As seen in Table 3.6, the SER accuracy (i.e., UA) drops with λ increases, whereas WER barely changes. It is likely because the smaller the ASR training weight, the longer time ASR reaches convergence so that the W2V2 representations can better be adapted to SER task. On the other hand, as it takes only two to three epochs for ASR training to converge as we noted during the

Table 3.5: Layer-wise analysis results. $\lambda = 0.1$.

Model	Hidden Output	UA
ASR-SER	layer 0 reps.	59.2%
	layer 1	60.4%
	layer 2	60.8%
	layer 3	61.5%
	layer 4	62.8%
	layer 5	63.5%
	layer 6	64.4%
	layer 7	64.0%
	layer 8	63.0%
	layer 9	62.5%
	layer 10	62.1%
	layer 11	55.3%
	layer 12	51.9%

Table 3.6: The impact of ASR training weight λ .

λ	UA	WER
0.1	68.4%	18.9%
0.3	67.2%	18.8%
0.5	66.6%	18.8%
0.7	66.0%	18.9%
0.9	64.7%	18.8%

experiments, λ barely impacts the ASR performance (i.e., WER). In contrast to H. Feng et al. (2020) which noted a large WER drop using a conventional encoder-decoder ASR model in joint ASR-SER training, wav2vec 2.0 demonstrated the potential for achieving satisfactory performance in both tasks through speech foundation models.

3.3.4 Summary

Following the experiments in Section 3.2, this section extends the analysis of wav2vec 2.0 by integrating its ASR training purpose into a joint training with SER. We compared the joint ASR-SER performance using different acoustic inputs and ASR outputs, as well as investigated the effectiveness of layer-wise hidden output and the impact of ASR training. The findings contribute to better understanding of the relationship between ASR and SER through speech foundation models.

3.4 Chapter Summary

This chapter presents a comprehensive investigation into the relationship between the speech foundation model wav2vec 2.0 and SER. Section 3.2 conducts a series of probing experiments on emotional speech using W2V2. Building on the findings from Section 3.2, Section 3.3 introduces a joint ASR-SER training framework that incorporates ASR outputs into SER and integrates paralinguistic features to address the paralinguistic shortage of wav2vec 2.0. A limitation of this chapter is its exclusive focus on the wav2vec 2.0 series. Although other self-supervised models may exhibit similar behavior, further investigation is worthwhile for comparison. Additionally, the supervised ASR model Whisper is expected to demonstrate distinct patterns. Nevertheless, the work presented in this chapter lays the groundwork for exploring ASR models for SER, providing a foundation for leveraging the strengths of ASR while mitigating its weaknesses.

The next chapter builds upon this one by shifting focus from acoustic to linguistic aspects, emphasizing the analysis of ASR lexical output.

Utilizing ASR Transcription for SER

4.1 Introduction

Recent work on SER consistently shows the benefits of incorporating both textual and acoustic features on benchmark corpora (Sebastian et al., 2019; Yoon et al., 2018). Although some researchers have also attempted to translate this approach into the wild, e.g., in-car voice systems (Y. Li, 2021a), it is still rare to see SER applications in our daily lives. One of the major reasons is that the majority of SER research uses human annotation, i.e., gold-standard manual transcriptions. In contrast, even for the ‘lab’ emotion corpora, transcriptions from a SOTA ASR system can result in high WERs. This means that very few of the findings obtained in the lab can be replicated in the wild.

To bridge this gap, it is necessary to exploit the imperfect textual data generated by ASR for emotion recognition. While previous studies proposed using ASR transcriptions for SER (Sahu et al., 2019b), the effect of ASR on emotion and vice-versa is not always clear. For example, a high WER on emotional speech is generally assumed to be a result of the distortion relative to neutral speech (Fernandez, 2004). Thus, even though ASR is a relatively well studied area, key questions about how it can be used in SER remain unresolved, such as 1) how confident can we be of ASR on different emotions? and 2) how is the SER performance affected by ASR errors? Thus, to make true progress in SER, we need to understand the interrelationship between ASR and SER.

In this section, we look at these fundamental yet long-standing issues by exploring the relationship between ASR results and emotional speech from the linguistic perspective, with the objective of pushing SER research closer to realistic use scenarios. Specifically, we first take a deep look into the word distribution of emotion corpora and how they are misrecognized in speech recognition. We analyze WER variation to see how ASR performance is affected by different emotions and how this relates to word-level confidence scores.

4.2 Exploring ASR Transcription of Emotional Speech

4.2.1 Experimental Setup

Datasets

We use three corpora: IEMOCAP, CMU-MOSI, and MELD to cover as many speech conditions as possible for generalizability. For IEMOCAP (IEM), we use four emotion classes: *angry*, *happy (+excited)*, *neutral*, and *sad*. For IEM and MELD, we removed utterances whose transcription is blank or whose audio file is too long due to mis-processing in corpora construction, bringing the total utterance numbers to 5 500 and 13 689, respectively.

ASR Models

For the ASR models, we use the Kaldi Librispeech ASR model¹ (KLA), a self-supervised model: wav2vec2-base-960h² (W2V2), a Conformer model³ (CONF) from ESPnet (Watanabe et al., 2018), and the medium model of Whisper⁴ (WHIS). KLA, W2V2, and CONF are pre-trained on Librispeech960, and WHIS on 680 000 h of multilingual and multitask data collected from the web. The ASR systems output whole words instead of characters. We use these four models for generalizability.

4.2.2 Experimental Results

In this section, we analyze the influence of emotion on ASR by conducting a series of experiments to identify sources of errors across corpora and word classes and quantifying the effect of WER on SER performance.

4.2.3 WER on Emotional Speech

Table 4.1: WER (%) of the ASR models on the emotion corpora.

	IEM	MOSI	MELD
KLA	36.8	40.9	58.5
W2V2	32.7	35.4	57.8
CONF	27.1	30.1	52.1
WHIS	12.3	17.3	34.8

1. <https://kaldi-asr.org/models/m13>

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

3. <https://github.com/espnet/espnet/blob/master/egs/librispeech/asr1/RESULTS.md>

4. <https://openai.com/research/whisper>

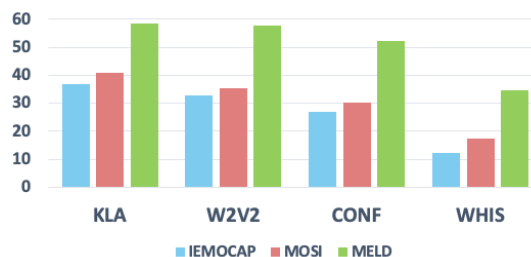


Figure 4.1: WER (%) of the ASR models on the emotion corpora.

We compute the WERs of the four ASR models on the three emotion corpora as the basis for analysis, as shown in Table 4.1 and Figure 4.1. Among the four models, WHIS shows the best performance on all the corpora, greatly outperforming the other three. We see large performance differences among corpora, even though they are all US or UK English emotional speech. This could be due to the different recording settings, the noise that occurred, or the vocabulary that was said.

4.2.4 Exploring ASR Errors Based on Part-of-Speech

In addition to the well-known acoustic factors that affect ASR (e.g., recording quality, noise), we would like to know if lexical factors, such as word distribution also affects WER. To investigate this, we analyze the words in each corpus based on two aspects: Part-of-Speech (PoS) tag and affective score. In particular, we explore whether words of a particular class or with strong affective coloring have a strong impact on ASR performance.

PoS. We select seven classes of PoS tags: *Noun*, *Verb*, *Adjective*, *Adverb*, *Wh-word*, *Function-word* from the Penn Treebank (Marcus, Santorini, & Marcinkiewicz, 1993) and *Stop-word* from the stopwords corpus of NLTK (Bird, 2006). The tagging was also conducted using NLTK.

Affective score. We use an affective words database that has nearly 14 thousand English words (Warriner et al., 2013) to refer the affective scores. Each word is rated in three affective dimensions: Valence (V), Arousal (A), and Dominance (D), ranging from 1 (weak) to 9 (strong). We divide the scores into three classes: low (1-3), medium (4-6), and high (7-9) according to their overall mean value, respectively.

We define three metrics to see how different classes of words affect the ASR performance. The Word Ratio (WR) is used to see the distribution of different classes of words in a corpus. The Error Ratio (ER) indicates if misrecognized words largely come from a specific class. The Class Error Rate (CR) measures how difficult words in a class are to recognize. We define them as follows.

$$\text{word ratio} = \frac{\text{word count per class}}{\text{total word count}} \quad (4.1)$$

$$\text{error ratio} = \frac{\text{word error count per class}}{\text{total word error count}} \quad (4.2)$$

$$\text{class error rate} = \frac{\text{word error count per class}}{\text{total word count per class}} \quad (4.3)$$

The WR⁵, ER, and CR based on PoS tag and affective score are shown in Table 4.2 and 4.3, respectively, with illustrative examples from IEMOCAP shown in Figure 4.2 and Figure 4.3. For brevity, we averaged the values over the four ASR models. As the ER values from the four ASR models are close and the CR values are proportionally distributed, taking the average does not affect the general findings.

Table 4.2: WR, ER, and CR (all in %) based on PoS tag.

	IEM			MOSI			MELD		
	WR	ER	CR	WR	ER	CR	WR	ER	CR
<i>Noun</i>	19.1	34.8	33.3	21.7	36.8	45.3	25.3	44.0	58.0
<i>Verb</i>	22.5	10.6	26.0	19.3	8.4	25.9	20.3	8.2	34.5
<i>Adj</i>	5.7	2.0	19.4	10.6	3.5	27.5	5.7	2.0	32.1
<i>Adv</i>	8.6	7.0	25.6	8.6	5.1	23.5	7.6	5.9	38.8
<i>Wh</i>	2.1	1.4	26.0	1.7	0.7	21.0	2.4	1.8	44.1
<i>Func</i>	19.2	12.1	20.1	25.1	15.6	24.1	17.5	10.2	30.3
<i>Stop</i>	50.7	33.7	25.7	50.0	30.0	24.8	46.8	29.5	33.7

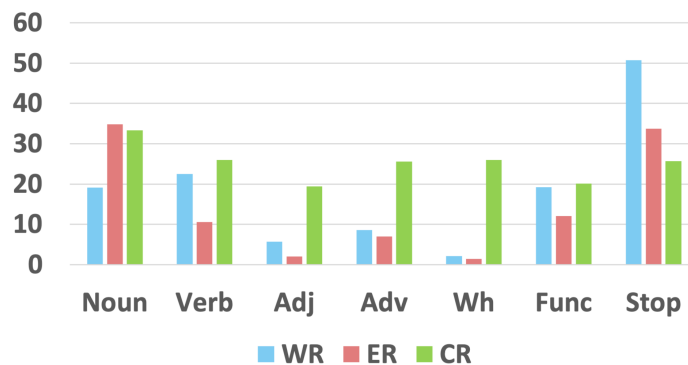


Figure 4.2: WR, ER, and CR (all in %) based on PoS tag (IEMOCAP).

Table 4.2 shows that ER generally increases with WR. This is reasonable, as the more words in one class, the more errors will be generated. However, this does not hold for all PoS classes: the WRs for *Noun*, *Verb* and *Func* are similar, but the ER is much lower for *Verbs*. That is, ASR performs better on verbs than expected based on verb frequency. On IEM, *Verb* has a higher WR than *Noun* and *Func* but much lower ER, and similarly for *Adj* and *Adv* on MOSI.

5. Note that *Stop* words can have multiple PoS tags as *Stop* is not in the Penn Treebank, making the summation of the word ratios greater than 100%.

This indicates that in ASR models, some classes of words are more difficult to recognize than others. The *Noun* class is the most difficult to recognize, showing the highest CRs across the corpora. However, other word classes do not have uniform patterns across the corpora. For instance, *Wh* words are the easiest to recognize on MOSI yet the second hardest on MELD.

The CRs of most word classes increase from IEM to MELD, which is in line with the change of WERs in Table 4.1. This is plausible as the word error count increases with the WER. Nevertheless, the rate of growth of CR varies from class to class. For example, *Noun* saw a 74% increase in CR (33.3 to 58.0) compared to 33% for *Verb* (26.0 to 34.5). Again, this reflects the fact that word classes have different ERs, perhaps because they have different tolerances for speech distortion.

4.2.5 Exploring ASR Errors Based on Affective Scores

Table 4.3: WR, ER, and CR (all in %) based on affective score.

	IEM			MOSI			MELD		
	WR	ER	CR	WR	ER	CR	WR	ER	CR
V_{low}	2.5	3.6	36.1	2.3	3.6	43.2	2.0	2.9	43.0
V_{mid}	10.8	13.8	29.6	9.3	11.7	33.9	9.8	14.2	37.7
V_{high}	13.9	16.2	26.3	15.6	18.9	34.3	12.2	16.2	35.5
A_{low}	15.0	17.4	24.8	12.3	13.8	30.9	13.2	17.3	35.0
A_{mid}	11.1	14.7	30.7	13.6	17.5	36.5	10.1	15.1	39.4
A_{high}	1.0	1.4	32.8	1.3	2.0	43.6	0.7	1.0	40.0
D_{low}	1.0	1.4	28.1	1.1	1.8	47.2	0.8	1.2	38.9
D_{mid}	16.7	20.5	27.7	15.6	20.0	36.2	15.0	20.9	36.9
D_{high}	9.4	11.5	27.4	10.6	11.5	29.2	8.1	11.2	36.7

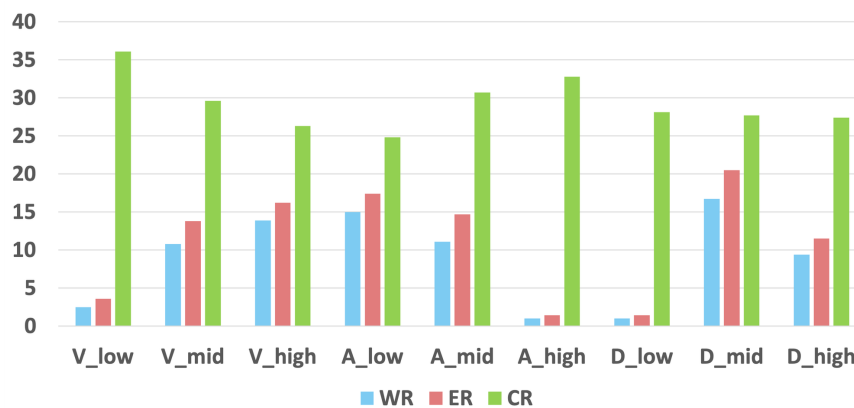


Figure 4.3: WR, ER, and CR (all in %) based on affective score (IEMOCAP).

Looking at the affective score-based distribution in Table 4.3 and Figure 4.3 and again see that, in general, the higher the WR, the higher the ER. However, in this case, we do not see classes with similar WRs but very different ERs. Nevertheless, differences still remain: CR decreases as valence and dominance scores increase, i.e., high valence and dominance words are better recognized. In contrast, we see that higher-arousal words are more difficult to recognize. Even so, unlike the classes based on PoS tags, the CRs have a relatively balanced distribution here, and their rates of increase with the WER are at the same level.

We see similar but not identical patterns when we look at the results for individual ASR models. We omit detailed results for brevity but note per-class differences in performance, e.g., *Nouns* are better recognized by KLA (32% ER) compared to CONF (35% ER) even though the latter has a lower overall WER. In general, no one model performed the best across all PoS or affective score classes.

4.2.6 Exploring WER with Utterance Length

One of the major factors that is likely to have an effect on ASR performance is utterance length. Generally, longer utterances are more likely to have more word errors compared to shorter ones because there are more opportunities for the ASR model to make errors (Sahu et al., 2019b). However, we hypothesize that longer utterances contain more contextual information that certain ASR models (e.g., W2V2) can use to compensate for the speech distortion.

To investigate this, we analyze how WER and WR change with utterance length. The four ASR models displayed the same trends, so Table 4.4 and Figure 4.4 shows results for W2V2 for brevity. We see that the shorter the utterances, the higher the WER, regardless of corpus. These support our hypothesis and additionally show that if a corpus contains many short utterances, ASR may not work well.

Table 4.4: WER and ratio (both in %) according to different utterance lengths (number of words N).

N	IEM		MOSI		MELD	
	Ratio	WER	Ratio	WER	Ratio	WER
≤10	59.6	46.4	53.9	40.6	71.5	73.3
11-20	22.9	32.2	32.2	34.5	23.5	48.6
21-30	10.8	26.3	9.5	31.9	4.4	42.1
≥30	6.7	25.0	4.4	32.6	0.6	38.8

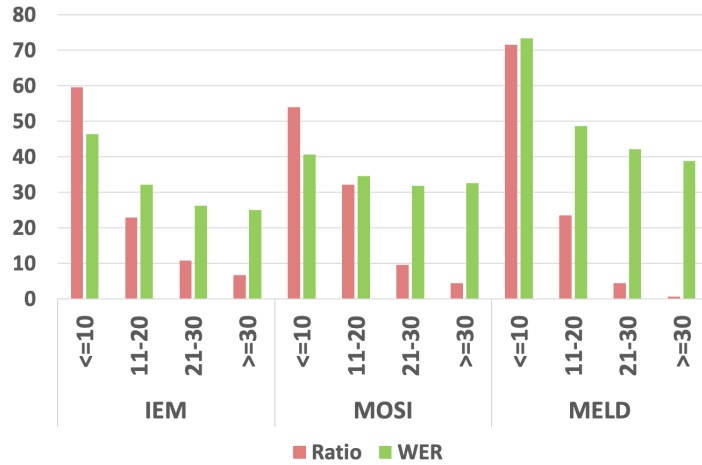


Figure 4.4: WER and ratio (both in %) according to different utterance lengths (number of words N).

4.2.7 Exploring ASR Performance on Different Emotions

It has long been argued that the acoustic characteristics of emotional speech, e.g., prosody variation, can deteriorate ASR performance (Goldwater et al., 2010). Word-level factors, however, are rarely studied. Based on the findings from the above experiments, we conducted a word-level analysis based on different emotions. Specifically, we look at the WER alongside the ratio of *Noun* and short utterance (word count less than 10), which we found to be major factors affecting ASR performance above.

As MOSI does not contain discrete emotion labels, we omit it in this exploration task. From Table 4.5, we can see the WER on Neutral is not the best but instead worse than most of the other emotions in both corpora. As neutral speech is the least emotionally “distorted”, the reason could be that the ratio of short utterances in Neutral is very high, which offsets the benefit of less emotional distortion. We did not see a large discrepancy in the ratio of *Noun* between emotions, although we found that a high ratio harms ASR performance in Section 4.2.4.

Table 4.5: WER, ratio of *Noun* words, and short utterances (all in %) according to different emotions on IEM and MELD.

<i>IEM</i>	Ang	Hap	Neu	Sad
WER	22.8	38.9	36.3	29.5
<i>Noun</i>	20.0	21.3	20.4	19.5
≤ 10	44.8	55.7	60.5	52.7

<i>MELD</i>	Ang	Dis	Fea	Joy	Neu	Sad	Sur
WER	52.7	53.4	58.6	59.9	58.3	52.9	65.3
Noun	26.2	27.6	26.4	29.7	26.9	25.2	24.5
≤10	64.1	60.1	62.1	72.4	72.8	62.0	82.8

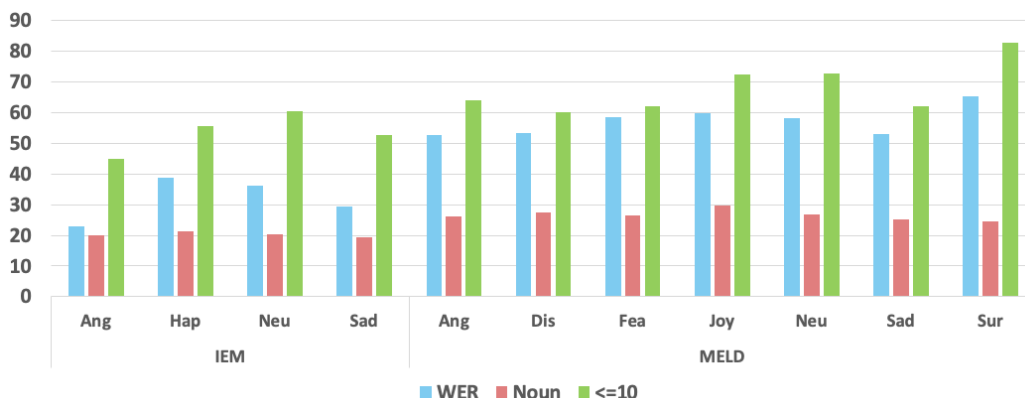


Figure 4.5: WER, ratio of Noun words, and short utterances (all in %) according to different emotions on IEM and MELD.

4.2.8 Exploring ASR Confidence Scores Across Emotions

Confidence scores can help identify words of high recognition quality for better use of ASR transcriptions. Santoso, Yamada, Makino, Ishizuka, and Hiramura (2021) and Santoso, Yamada, Ishizuka, Hashimoto, and Makino (2022) adjusted the attention weights for words using their confidence scores. Pan et al. (2020) proposed removing words with low confidence and selecting words with the highest confidence from multiple ASR hypotheses (Pan et al., 2021). However, the measurement and use of confidence scores is a complicated task due to many reasons such as out-of-domain words and overconfidence (Q. Li et al., 2022; Qiu et al., 2021). Hence, we analyze how confidence varies with emotion by exploring the word confidence scores. We compared two ASR models: CONF and KLA (as the confidence scores produced by the two models are more accessible), and also compared the confidence scores generated by KLA with and without calibration. The calibration was applied by using the official procedure in Kaldi⁶. Since we noticed similar patterns of confidence in all three corpora, we only show the results of IEM for brevity.

6. https://github.com/kaldi-asr/kaldi/blob/master/egs/tedlium/s5/local/confidence_calibration.sh

In Table 4.6, we observe a great difference in confidence scores between ASR models. The scores for correctly recognized words are similar between CONF and KLA, but for incorrectly recognized words, the score is 0.70 in KLA compared to 0.87 in CONF. Overall, KLA scores are generally lower than those of CONF. Additionally, in the emotions Angry, Neutral, and Sad, confidence scores increase with utterance length: the longer the utterance, the higher the confidence, which aligns with the WER trend shown in Table 4.4. However, this pattern does not appear in Angry.

Table 4.6: Distribution of confidence scores of the ASR models. (The CONF column is illustrated in Figure 4.6. The last two columns are illustrated in Figure 4.7).

	N	CONF	KLA	KLA (calibrated)
Correct		0.97	0.95	0.62
Incorrect		0.87	0.70	0.53
Ang	≤ 10	0.96	0.91	0.62
	11-20	0.96	0.91	0.61
	21-30	0.96	0.90	0.61
	≥ 30	0.96	0.90	0.61
Hap	≤ 10	0.92	0.80	0.57
	11-20	0.94	0.83	0.58
	21-30	0.95	0.86	0.59
	≥ 30	0.94	0.85	0.58
Neu	≤ 10	0.91	0.81	0.58
	11-20	0.94	0.86	0.59
	21-30	0.95	0.88	0.59
	≥ 30	0.96	0.89	0.60
Sad	≤ 10	0.92	0.81	0.58
	11-20	0.95	0.87	0.59
	21-30	0.96	0.89	0.60
	≥ 30	0.97	0.92	0.61

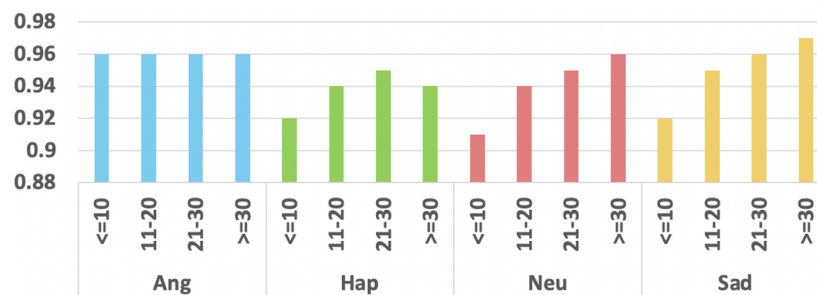


Figure 4.6: Distribution of confidence scores of CONF model.

By comparing the confidence scores with and without calibration for KLA, we observe that the gap in confidence between correctly and incorrectly recognized words becomes smaller. Similarly, the influence of utterance length nearly disappears, indicating that emotions are treated more equally after calibration. Calibration also mitigates the overconfidence commonly

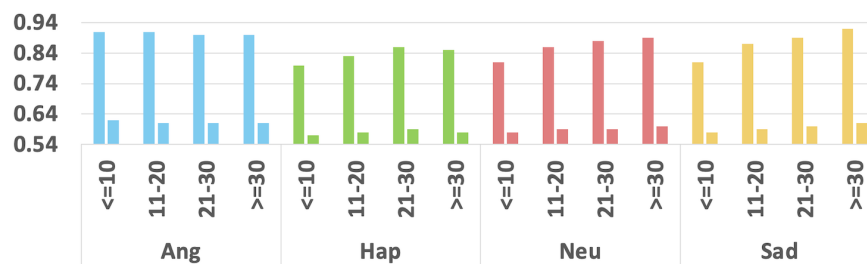


Figure 4.7: Distribution of confidence scores of KLA model w/o and w/ calibration (left and right per utterance length).

exhibited by end-to-end ASR models (Q. Li et al., 2021). The overall confidence for each emotion corresponds to their respective WERs in Table 4.5, with ASR performance being the best on Angry and the worst on Happy. These findings highlight that confidence scores vary with emotion, utterance length, and ASR model. Therefore, these variations need to be carefully considered and confidence scores properly calibrated when they are used to adjust word contributions in SER.

4.2.9 Summary

In this section, we addressed a long-existing yet understudied issue: the relationship between ASR lexical output and emotion. We demonstrated that word distribution in a corpus is a key factor that affects ASR performance and that differences exist between emotions. The frequencies of different PoS tags, affective scores, and utterance lengths may have a large impact on ASR performance. For example, the more the Noun words and short utterances, the higher the WER, regardless of the speaking style (speech collected in the lab, from monologues or TV shows). Moreover, different ASR models behave similarly in these trends, but with minor differences: some classes of words are easier to recognize for one model but maybe harder for another. Moreover, the confidence scores generated by different models and with calibration may show large differences, so the use of confidence scores should be carefully examined. We expect these findings to provide a valuable foundation for future SER research using ASR transcription, such as knowledge integration for LLM prompting.

4.3 Impact of ASR Transcription Quality on SER

After examining ASR transcription of emotional speech, we now investigate how the quality of ASR transcription affects SER. Over the past several years, significant progress has been achieved in SER through the utilization of ASR transcriptions, specifically, in chronological order:

Yoon et al. (2018) proposed a deep dual recurrent encoder model that simultaneously utilizes speech signal and transcriptions from the Google Cloud Speech API. Sahu et al. (2019b) utilized two commercial ASR systems to generate transcriptions for bimodal SER (speech + text), resulting in a relative loss of unweighted accuracy compared to ground-truth transcription. H. Li, Ding, Wu, and Liu (2020) introduced a temporal alignment mean-max pooling mechanism to capture subtle and fine-grained emotions in utterances, alongside a cross-modality excitement module for sample-specific adjustments on embeddings.

Santoso et al. (2021) proposed using a confidence measure to adjust the importance weights in ASR transcription based on the likelihood of a speech recognition error in each word, effectively mitigating the effects of ASR errors on SER performance. W. Wu et al. (2021) introduced a dual-branches model for ASR error robustness, with a time-synchronous branch combining speech and text modalities and a time-asynchronous branch integrating sentence text embeddings from context utterances. Shon, Brusco, Pan, Han, and Watanabe (2021) generated pseudo labels on ASR transcription for semi-supervised speech sentiment analysis.

While these studies have highlighted the effectiveness of integrating text features from ASR transcription, there is still a lack of understanding regarding how WER and fusion techniques impact SER. Therefore, we undertake a benchmark study utilizing diverse ASR models and fusion techniques, conducting SER on various emotion corpora to get a clearer picture of the effect of transcription errors on SER.

4.3.1 Experimental Setup

ASR models

We adopt the following 11 models as they are widely used in the speech area and can provide varying WERs.

- *wav2vec2-base-{100h,960h}*
- *wav2vec2-large-960h*
- *wav2vec2-large-960h-lv60-self*
- *HuBERT-large-ls960-ft*
- *WavLM-libri-clean-100h-base-plus*
- *Whisper-{tiny, base, small, medium, large-v2}.en*

Datasets

To ensure generalizability, we utilize three corpora: IEMOCAP, CMU-MOSI, and MSP-Podcast, to include diverse speech conditions and various evaluation metrics for both discrete and continuous emotions.

For IEMOCAP, we focus on four emotion classes: *angry*, *happy (including excited)*, *neutral*, and *sad*. We exclude utterances with blank transcription from this corpus, resulting in a total of 5 500 utterances. CMU-MOSI comprises 2 199 monologue video clips sourced from YouTube, each annotated with sentiment scores ranging from -3 to 3. For MSP-Podcast, we use its Release 1.11 version and evaluate its performance on the Test1 set. We exclude utterances without an emotion label, resulting in a total of 104 663 utterances.

Following the literature (Hazarika, Zimmermann, & Poria, 2020; Tsai, Bai, et al., 2019), we compute **Acc4** (unweighted four-class accuracy) for IEMOCAP, Concordance Correlation Coefficient (**CCC**) for MSP-Podcast, and **Acc2** (binary accuracy), **Acc7** (seven-class accuracy), and Mean Absolute Error (**MAE**) for CMU-MOSI.

SER model

We employ RoBERTa-base as the text encoder, as it is widely used for text-based SER (J. Luo, Phan, & Reiss, 2023; Siriwardhana et al., 2020). Given that transcription generated by models other than *Whisper* lack punctuation, we remove punctuation from the transcription of *Whisper* models to ensure a fair comparison. All letters are lowercased for consistency. A backbone SER model is built for all corpora. Since our goal is not to achieve state-of-the-art performance, the model simply comprises two dense layers: the first encodes RoBERTa output of dimension 768 into hidden states of dimension 128, and the second further encodes it into a dimension of 16. We use ReLU as the activation function between the dense layers. In the case of IEMOCAP, we apply one output layer with Softmax activation for classification. For MSP-Podcast and CMU-MOSI, which involve regression tasks, no final output activation is applied. We set the learning rate as 5.0×10^{-4} for IEMOCAP and CMU-MOSI, and 1.0×10^{-4} for MSP-Podcast, using the AdamW optimizer. The weight decay is set as 1.0×10^{-5} , and the batch size is 64. For training, we employ five-fold cross-validation on IEMOCAP (100 epochs) and follow the official training/validation/testing sets on CMU-MOSI (100 epochs) and MSP-Podcast (30 epochs). The random seeds are kept consistent across all experiments. Note that the SER model is trained and tested on the same transcription source, rather than being trained on ground truth and then tested on ASR transcription.

4.3.2 Benchmarking SER Across WERs with Text Only

We present the WERs and corresponding SER performance based on transcription from different ASR models. From Table 4.7, we observe that:

1. SER performance generally decreases as WER increases. On IEMOCAP and MOSI, there is nearly a 10% accuracy decrease with WERs around 40%, regardless of Acc2, Acc4, or Acc7. However, exceptions exist. For example, on IEMOCAP, *Whisper-tiny* produces higher accuracy than its neighboring models, even though the WERs are similar. Additionally, *W2V960-large-self* yields worse accuracy compared to its neighboring models, despite having a relatively lower WER compared to others. The same phenomenon can also be found with CMU-MOSI: the Acc2 of *W2V100* is higher than that of *WavLM plus*, and the Acc2 of *Whisper-large* is lower than that of *Whisper-medium*. This might be due to certain words being misrecognized as words that have little effect on or even positively contribute to their ground-truth emotion labels.
2. SER is robust to relatively low WER, and in some cases, it is even better with ASR errors. From IEMOCAP, it is observed that a WER of approximately 12% has minimal impact on SER performance compared to ground-truth transcription. Moreover, in CMU-MOSI, ASR errors can potentially enhance SER: transcription generated by specific *Whisper* models outperform those based on ground truth. Prior research has shown that sentiment analysis remains robust against ASR errors (Tokuhisa, Inui, & Matsumoto, 2008), as both positive and negative sentiments can consist of a range of emotional states and therefore require less granularity. Additionally, we find that in certain cases (e.g., relatively low WER), ASR errors do not diminish SER performance. Recent research on speech-based dementia detection also indicates that ASR errors can offer valuable insights for dementia classification (C. Li, Xu, Cohen, & Pakhomov, 2024), possibly due to specific types of errors (e.g., repetitions, disfluencies) that may reveal indicators of dementia.
3. Different metrics have different sensitivities to WER. On CMU-MOSI, Acc7 shows a distinct pattern compared to Acc2 and MAE. Acc2 and MAE demonstrate consistent and smooth variations, whereas Acc7 appears random and lacks a discernible pattern. Interestingly, in nearly half of the cases (all from *Whisper*), automatic transcription outperforms the ground-truth transcripts in terms of Acc7. This discrepancy may stem from the mismatch between the regression model used during training (since MOSI is labeled with continuous values) and the classification metric applied to Acc7, where predicted and ground-truth values are grounded (i.e, the Acc7 is based on grounded regression score). This mismatch could potentially blur accuracy and render SER insensitive to WER.

Table 4.7: Benchmarking SER performance based on ASR transcription across corpora. \uparrow : higher the better. \downarrow : lower the better. **Red bold**: better performance than the ground truth.

IEMOCAP				
ASR Model	WER \downarrow	Acc4 \uparrow		
w2v100	40.31	52.80		
wavlm plus	39.14	52.84		
w2v960	32.37	55.16		
w2v960 large	27.98	57.05		
whisper tiny	26.97	59.54		
hubert large	23.80	58.65		
w2v960 large self	21.90	57.35		
whisper base	20.27	60.77		
whisper small	14.44	61.89		
whisper medium	12.65	62.32		
whisper large	12.31	62.95		
ground truth	0.00	63.20		

CMU-MOSI				
ASR Model	WER \downarrow	Acc2 \uparrow	Acc7 \uparrow	MAE \downarrow
w2v100	41.83	72.41	27.70	1.1216
wavlm plus	41.62	69.82	25.66	1.1351
w2v960	35.93	73.63	29.88	1.0855
w2v960 large	31.65	75.61	33.82	1.0279
hubert large	28.89	76.83	29.15	0.9942
w2v960 large self	27.40	78.81	31.92	0.9876
whisper tiny	27.31	79.27	35.13	0.9545
whisper base	22.72	80.18	34.55	0.9246
whisper small	21.81	82.01	37.03	0.8998
whisper medium	18.96	83.23	36.59	0.8668
whisper large	18.69	82.01	36.88	0.8561
ground truth	0.00	82.47	34.26	0.8974

MSP-Podcast				
ASR Model	WER \downarrow	CCC_v \uparrow	CCC_a \uparrow	CCC_d \uparrow
w2v100	40.14	0.465	0.331	0.299
wavlm plus	37.67	0.467	0.322	0.286
w2v960	33.84	0.479	0.316	0.288
w2v960 large	30.85	0.489	0.318	0.289
hubert large	27.74	0.504	0.330	0.298
w2v960 large self	25.12	0.494	0.311	0.286
whisper tiny	24.63	0.503	0.313	0.286
whisper base	19.35	0.521	0.315	0.289
whisper small	15.60	0.525	0.318	0.290
whisper medium	15.14	0.525	0.315	0.291
whisper large	15.02	0.526	0.316	0.293
ground truth	0.00	0.524	0.321	0.301

4. Different emotion dimensions exhibit varying degrees of robustness to ASR errors. From MSP-Podcast, it is evident that valence, arousal, and dominance exhibit distinct patterns. Firstly, the CCC of valence mirrors the pattern observed in Acc4 for IEMOCAP and Acc2 for CMU-MOSI, suggesting that valence shares similarities with categorical emotion in terms of robustness to ASR errors. Given that valence is conceptually similar to sentiment (indicating positivity or negativity), this alignment is plausible. Secondly, arousal and dominance do not exhibit a clear correlation with WERs. Our observations are consistent with the conventional understanding that valence is more influenced by textual content, whereas arousal and dominance are more influenced by audio cues (Y. Li, Ishi, et al., 2019; Schuller, Vlasenko, Eyben, Rigoll, & Wendemuth, 2009; Wagner et al., 2023).

Indeed, when we replaced text features with middle-layer speech features from *W2V960-base* in the SER model, we obtained CCC_v, CCC_a, and CCC_d values of 0.531, 0.635, and 0.558, respectively, further validating our observations.

4.3.3 Benchmarking SER Across WERs with Bimodal Fusion

We integrate speech features to investigate the robustness of existing fusion techniques in handling ASR errors in real-world scenarios. Since speech is not the primary focus of this study, we simply use the middle-layer speech features from *W2V960-base*. For fusion techniques, we employ the following approaches:

- **Early fusion**: text and speech features are concatenated at the embedding level.
- **Late fusion**: text and speech features are learned independently by their respective models, and the final decision is determined based on their outputs.
- **Cross-attention fusion**: text and speech features are attended to each other via attention mechanism and then concatenated.
- **Tensor fusion** (Zadeh, Chen, Poria, Cambria, & Morency, 2017): unimodal information and bimodal interactions are learned explicitly and then aggregated.
- **Non-local gate-based (NL-gate) fusion** (X. Wang, Girshick, Gupta, & He, 2018): NL-gate is incorporated with multiple directions and at multiple positions with query-key-value within the attention mechanism.
- **Modality-invariant and -specific fusion (MISA)** (Hazarika et al., 2020): combining both modality-invariant and modality-specific features.

Table 4.8: Benchmarking SER performance based on ASR transcription with fusion techniques. Maximum diff: maximum difference between the performance of ground truth and that of ASR transcription. **Red bold:** better performance than the ground truth.

ASR Model	WER	Text	IEMOCAP (Acc4)					MISA
			Early	Late	Attn	Tensor	NL-gate	
w2v100	40.31	52.80	69.12	66.37	68.46	70.49	67.26	70.15
wavlm plus	39.14	52.84	68.84	66.88	69.10	70.45	67.00	70.32
w2v960	32.37	55.16	69.43	67.69	69.47	70.98	68.51	71.43
w2v960 large	27.98	57.05	69.85	68.15	69.65	71.47	68.62	71.87
whisper tiny	26.97	59.54	71.20	68.90	70.79	72.73	70.02	73.05
hubert large	23.80	58.65	70.62	68.64	70.73	72.32	69.06	72.66
w2v960 large self	21.90	57.35	70.85	68.44	70.61	72.29	69.08	72.69
whisper base	20.27	60.77	72.63	69.60	71.54	73.73	70.96	73.60
whisper small	14.44	61.89	72.61	70.14	72.50	74.53	71.10	73.98
whisper medium	12.65	62.32	72.74	70.28	72.96	74.63	70.97	74.36
whisper large	12.31	62.95	73.01	70.68	73.03	74.66	71.68	74.51
ground truth	0.00	63.20	72.68	70.52	72.99	74.16	71.45	74.32
maximum diff		10.40	3.84	4.15	3.51	3.67	4.44	4.00
MSP-Podcast (Average value of CCC_v, CCC_a, and CCC_d)								
ASR Model	WER	Text	Early	Late	Attn	Tensor	NL-gate	MISA
w2v100	40.14	0.365	0.603	0.532	0.594	0.481	0.579	0.585
wavlm plus	37.67	0.358	0.603	0.528	0.599	0.495	0.583	0.586
w2v960	33.84	0.361	0.603	0.528	0.596	0.482	0.581	0.587
w2v960 large	30.85	0.365	0.608	0.532	0.602	0.498	0.587	0.593
whisper tiny	27.74	0.377	0.613	0.536	0.605	0.494	0.590	0.598
hubert large	25.12	0.363	0.610	0.543	0.605	0.496	0.585	0.593
w2v960 large self	24.63	0.367	0.607	0.538	0.610	0.492	0.582	0.593
whisper base	19.35	0.375	0.613	0.537	0.611	0.497	0.592	0.599
whisper small	15.60	0.378	0.616	0.544	0.612	0.501	0.594	0.600
whisper medium	15.14	0.377	0.615	0.548	0.612	0.496	0.593	0.601
whisper large	15.02	0.378	0.616	0.549	0.613	0.499	0.594	0.601
ground truth	0.00	0.382	0.613	0.545	0.611	0.498	0.595	0.598
maximum diff		0.024	0.010	0.017	0.017	0.016	0.016	0.014

For fairness, we keep the backbone SER model used in the previous section unchanged, modifying only the input dimension of the first dense layer to match the output dimension of the hidden states from each fusion model. Additionally, we modify *Tensor Fusion* and *MISA* to receive bimodal inputs, as they were originally designed for trimodal inputs. We present results for IEMOCAP and MSP-Podcast to cover both categorical and dimensional emotions. Regarding CMU-MOSI, integrating speech did not largely improve performance, as text alone already yields strong results. Results are presented in Table 4.8.

It can be seen that:

1. Fusing speech features largely mitigates the negative impact of increasing WER. The decrease in Acc4 based on WER reaches 10% without fusion on IEMOCAP, but only 4% with fusion. Moreover, with most fusion techniques, the SER performance is even better on transcription with relatively low WER than on ground truth. These findings confirm the benefits of bimodal SER in real applications, which can be achieved by using a pre-trained ASR model as a base for jointly generating transcription and speech features (Cai et al., 2021; Y. Li, Bell, & Lai, 2022).
2. There is no optimal WER-robust fusion approach. Due to different amounts of encoded information, the performances of fusion techniques are not directly comparable. For instance, *Tensor Fusion* and *MISA* encode more information by incorporating both unimodal and bimodal features, or both modality-invariant and -specific features. However, their effectiveness varies across different corpora. For example, they perform well on IEMOCAP (categorical) but poorly on MSP-Podcast (dimensional), indicating large inconsistency. Thus, it is difficult to identify a single best-performing technique that is universally effective in addressing ASR errors. Nevertheless, the maximum difference between the performance of ground truth and that of ASR transcription is still indicative of the robustness of the fusion techniques. We observe that although *cross-attention fusion* may not produce the best performance, it yields a relatively small maximum difference, suggesting that it is less affected by WER. This observation is reasonable since cross-attention captures the relatedness between bimodal inputs, which changes dynamically once the words in text have changed, ensuring that the relatedness is least affected by ASR errors.

4.3.4 Summary

We conducted a benchmark of SER performance using ASR transcripts with varying WERs and explored mainstream fusion techniques to assess the impact of ASR performance on SER. Our findings revealed several novel insights: **1)** SER can tolerate relatively low WERs, especially in real-life speech scenarios. **2)** Bimodal SER with transcripts containing approximately 10% errors may not perform worse than those with ground-truth text, particularly with powerful Whisper models. However, further analysis is necessary to understand the nature and locations of ASR errors, as well as the mechanisms underlying fusion techniques. Based on this work, we have launched a challenge aimed at addressing ASR errors for text-based SER (C.-H. H. Yang et al., 2024).

4.4 Chapter Summary

This chapter presents a comprehensive study of ASR transcription of emotional speech. Section 4.2 investigates the ASR transcription through word-level analysis to understand how different emotions affect specific words. Section 4.3 conducts a benchmark study on the impact of WER in SER, examining how WER and fusion techniques influence both classification and regression performance. This chapter contributes to a deeper understanding of the relationship between ASR's lexical output and emotional speech, aiming to leverage ASR transcription in SER.

The next chapter builds on this foundation by introducing novel ASR error correction techniques to improve ASR transcription quality for enhanced SER performance.

Correcting ASR Errors for SER

5.1 Introduction

Building on the findings from the previous chapter, we now focus on addressing the challenges posed by ASR transcription errors in the context of SER. In this chapter, we explore ASR Error Correction (AEC), a post-processing approach, to enable more reliable use of ASR transcription for SER. Specifically, we explore emerging Large Language Models (LLMs) in Section 5.2 and traditional Sequence-to-Sequence (S2S) in Section 5.3 for the best application, hoping for joint AEC and SER.

5.1.1 LLM-Based Approaches

Recent studies have suggested that LLMs have the ability to reason about emotional content (Sap, Rashkin, Chen, Le Bras, & Choi, 2019). This finding has encouraged researchers to further explore the “emotional intelligence” (e.g., emotion recognition, interpretation, and understanding) of LLMs. For example, (X. Wang, Li, Yin, Wu, & Liu, 2023) developed a psychometric assessment focusing on emotion understanding to compare the emotional intelligence of LLMs and humans. They found that most LLMs achieved above-average Emotional Quotient (EQ) scores, with GPT-4 surpassing 89% of human participants with an EQ of 117.

Therefore, the use of LLMs in text-based emotion recognition has emerged as a resource-efficient and effort-saving alternative to human annotators and traditional emotion classifiers for two main reasons: **1)** Emotion recognition requires substantial human effort. Typically, multiple annotators are needed for each sample to reach a majority vote, ensuring accurate assessment. Although platforms like Amazon Mechanical Turk provide a relatively efficient solution, concerns persist regarding privacy leaks, subjective bias, and the reliability of annotations (T. Feng & Narayanan, 2024). **2)** Despite the advancements of state-of-the-art deep learning technologies, training an emotion classifier involves multiple steps, including feature extraction and model training, which require careful consideration of which features, models, and algorithms to use.

To this end, researchers have recently started exploring and utilizing LLMs for emotion annotation and recognition. These works include various approaches, such as using multiple-step prompting with LLMs (Hama, Otsuka, & Ishii, 2024), examining prompt sensitivity (Amin & Schuller, 2024), integrating outputs from multiple LLMs (T. Feng & Narayanan, 2024), and incorporating acoustic information (Santoso, Ishizuka, & Hashimoto, 2024). Despite these efforts, understanding of the usage, efficacy, and reliability of LLM-based approaches remains limited, especially on ASR transcription. Therefore, building upon existing literature, we conduct various experiments on LLM-based emotion recognition to investigate effective prompting practices.

5.1.2 S2S-Based Approaches

While S2S models have been established for AEC (Dutta et al., 2022; Hrinchuk, Popova, & Ginsburg, 2020; Y. Jiang & Poellabauer, 2021), research on Low-Resource Out-of-Domain (LROOD) scenarios is limited. Many previous studies were performed on the same large corpus without considering the LROOD problem (Tanaka et al., 2021; S. Zhang et al., 2021), leaving challenges remain such as determining effective Pre-Training (PT) and Fine-Tuning (FT) strategies with LROOD data. None of the prior work has considered the characteristics of the ASR models that are the source of transcript generation. Moreover, although some studies have used data augmentation to produce more erroneous sentences for LROOD downstream corpora for AEC training (Dutta et al., 2022; Ghosh et al., 2024), we argue that such arbitrary augmentation is unreliable because the error patterns of the augmented data differ from the original ASR errors. We hypothesize that different ASR models may produce distinct patterns of ASR errors due to their respective training process (e.g., training data, algorithms, targets), which requires that the AEC model be trained for corresponding ASR domain errors.

Additionally, acoustic information has proven useful for crossmodal AEC (B. Lin & Wang, 2023; Radhakrishnan et al., 2023), but it is not always possible to acquire audio sources for the PT stage (e.g., due to privacy or other ethical issues). Therefore, determining how to better incorporate audio features and which acoustic features are useful remains an open question, considering high-WER speech usually contains low-quality audio that can introduce distortions into the crossmodal training. Finally, very few studies have applied corrected ASR transcripts to downstream tasks to evaluate AEC extrinsically.

To improve ASR transcriptions for SER, we explore both LLM-based and S2S approaches in this chapter. For the LLM-based approach, we first investigate the use of LLMs for emotion recognition on ASR transcriptions, examining their effectiveness across varying levels of WER. We then introduce a novel prompting framework, Revise-Reason-Recognize (R3), which integrates AEC and reasoning with emotion-specific prompts to identify emotions from ASR transcriptions.

For the S2S-based approach, we propose a crossmodal method for AEC and conduct a thorough analysis comparing AEC performance with and without PT or FT on LROOD scenarios. Additionally, we compare transcriptions from different ASR models with similar WERs, enhance crossmodal AEC by incorporating speech representations only during the FT stage, and perform SER using transcripts corrected by our proposed AEC method.

5.2 LLM-Based Emotion Recognition on ASR Transcription

Interest in LLM-based emotion recognition has surged recently, with the availability of pre-trained models. Studies in this area have explored a variety of approaches: Gong et al. (2023) inferred emotion labels using three different prompting approaches: text generation, mask filling, and textual entailment, employing a fine-grained emotion taxonomy. T. Feng and Narayanan (2024) proposed an ensemble approach that integrates outputs from multiple LLMs, leveraging a Mixture of Experts (MoE) reasoning model. They trained emotion classifiers using MoE-generated emotion labels from both ground-truth and ASR transcriptions, and tested these classifiers on ground-truth labels, demonstrating comparable performance in emotion classification. Hama et al. (2024) employed a multi-step prompting technique with few training samples for text emotion recognition. Santoso et al. (2024) and Latif, Usama, Malik, and Schuller (2023) incorporated textual acoustic feature descriptors into prompts. Z. Zhang et al. (2024) investigated several approaches, including in-context learning, few-shot learning, accuracy, generalization, and explanation. Amin and Schuller (2024) examined various prompting techniques, including chain-of-thought, role-play, and their variations.

Despite these advancements, we argue that each approach has its limitations, leaving several concerns unaddressed. For example, Z. Zhang et al. (2024) did not study prompting, Hama et al. (2024) only proposed a multi-step prompting without further exploration. The prompting techniques tested by Amin and Schuller (2024) were not specifically designed for emotion. Santoso et al. (2024) and Latif et al. (2023) only considered basic acoustic features, without incorporating more emotion-related properties. Furthermore, both T. Feng and Narayanan (2024) and Gong et al. (2023) used ASR transcriptions generated by *Whisper*, whose output has been shown to be robust in emotion recognition even with ASR errors, as demonstrated in Chapter 4.3. However, this setting is not ideal for LLMs handling challenging ASR transcriptions in real-world emotion applications.

5.2.1 Methodology

To address the issues above, we **1)** Develop prompts that incorporate emotion-specific knowledge; **2)** Incorporate AEC to refine transcriptions for robust emotion recognition; and **3)** Explore LLM training schemes to further improve the performance.

Prompting with Emotion-Specific Knowledge

In light of the relationship between emotion and relevant disciplines, we extract useful knowledge from acoustics, linguistics, and psychology, to develop emotion-specific prompts for emotion recognition. The prompts are presented in Figure 5.1.

Acoustic information plays a crucial role in distinguishing speech emotions. Features like energy, pitch, and speaking rate have proven useful when been incorporated into prompts as textual descriptors (Santoso et al., 2024). Similarly, gender information, which is highly correlated with pitch, has also been shown to be useful (Latif et al., 2023; Y. Li, Zhao, et al., 2019). However, these features are insufficient to fully describe fine-grained differences in emotions beyond the Big Four. Hence, *we hypothesize that additional acoustic features can enhance LLMs' emotion recognition ability. We propose including pitch range, jitter, and shimmer to incorporate mid-level prosody (between frame-level and utterance-level) and voice quality.* The acoustic features are extracted using openSMILE ¹ and Mid-Level Prosodic Feature Toolkit ² calculated based on the threshold. The thresholds are taken from the lowest 30% and the highest 30% of the ordered numerical value of the acoustic feature. Each acoustic feature is classified as low, mid, or high.

Linguistic structure is essential for understanding emotion in text. For example, Section 4.2 investigated the impact of part-of-speech, affective score, and utterance length on emotion. However, to our knowledge, only one work, Singh, Caragea, and Li (2024) has utilized linguistics via identifying emotion triggers (i.e., words that elicit the emotion) when prompting LLMs for emotion prediction. Thus, *we hypothesize that LLM-based emotion recognition can benefit from more linguistic knowledge as LLM processing is inherently text-based. We propose including ASR-emotion relationships among emotion category, WER, and utterance length as outlined in Section 4.2.*

Psychological theories, such as self-monitoring, social cognitive theory, and cognitive emotion regulation have proven effective in improved LLMs' performance across various tasks (C. Li et al., 2023; X. Wang, Li, Chang, Wang, & Wu, 2024). Therefore, *we hypothesize that LLMs' emotion recognition ability can resonate with their emotional intelligence and thus be enhanced. We propose incorporating positive and negative stimuli from (C. Li et al., 2023; X. Wang et al., 2024), as well as create our novel competitive stimuli.*

1. <https://audeering.github.io/opensmile-python/>

2. <https://www.cs.utep.edu/nigel/midlevel/>

Baseline	No Reasoning	Reasoning
	Do not show your explanation.	Please give your explanation.
Acoustics	Gender-Informed F0	Paralinguistic Features
	The sentence is spoken by a female speaker with F0 value of xxx.	Loudness: (low/medium/high) Pitch: (low/medium/high) Pitch range: (narrow/medium/wide) Speaking rate: (slow/medium/fast) Jitter: (low/medium/high) Shimmer (low/medium/high)
Linguistics	Trigger Words	ASR Relationship
	Emotion triggers are the words that cause the emotion of the text sentence. You need to find the emotion triggers.	1. Short sentences (i.e., fewer than 10 words) usually result in a high WER. 2. Neutral has the highest ratio of short sentences, followed by happy, sad, and angry. 3. Happy sentences have the highest WER, followed by neutral, sad, and angry.
Psychology	Positive/Negative Stimuli	Competitive Stimuli
	Pos: This is very important to my career. You'd better be sure. Neg: Perhaps this task is just beyond your skill set.	GPT4 is doing very well on this. Can you do better?

Figure 5.1: Emotion-specific prompts used in this work.

Emotion Recognition with ASR Error Correction

Traditional emotion recognition often struggles with imperfect text as we have shown in Section 4.3. *We argue that it is more challenging to prompt LLMs for emotion recognition on ASR transcription compared to human transcription, due to the presence of word errors.* Therefore, we propose the R3 prompting pipeline to perform emotion recognition with AEC and reasoning on ASR transcriptions. The R3 pipeline involves three steps: **Revise**, where ASR errors are corrected based on N-best hypotheses; **Reason**, where the LLMs self-explain based on the corrected transcriptions and emotion-specific knowledge; and **Recognize**, where the emotion is recognized. To incorporate AEC into our prompts, we follow an AEC-specific Alpaca prompt (C.-H. H. Yang et al., 2023), which uses the “*You are an ASR error corrector*” instruction, guiding the LLMs to perform error correction. As LLMs have proven their ability in both AEC and emotion recognition (C.-H. H. Yang et al., 2024), this format is expected to facilitate seamless integration with our emotion prompting, instructing the LLMs to function simultaneously as both an ASR error corrector and an emotion recognizer.

Exploring LLM Training Schemes

To understand how LLM training schemes contribute to emotion recognition, we explore context-aware learning, in-context learning, and instruction tuning. For **context-aware learning**, we organize the sentences in the conversation order and compare different context windows (i.e., the number of sentences preceding the sentence to be recognized). For **in-context learning**, we test and compare several few-shot cases. For **instruction tuning**, we apply Parameter-Efficient Fine-Tuning (PEFT) using LoRA (E. J. Hu et al., 2021). PEFT using LoRA is a method for adapting large pre-trained models to new tasks without updating all the model parameters. Instead, it introduces a small number of trainable parameters into the model, significantly reducing the computational cost and memory usage of fine-tuning. We set learning rate, weight decay, and epoch to 1.0×10^{-4} , 1.0×10^{-5} , and 5, respectively, and keep the LoRA configuration at its default settings, after several times of model tuning.

5.2.2 Experimental Setup

For the **datasets**, we use IEMOCAP and the Test1 set of MSP-Podcast. We combine *excited* with *happy* and use the Big Four classes for IEMOCAP. For the **ASR models**, we adopt the following ten to generate diverse transcripts to form 10-best ASR hypotheses, following the model choice in Section 4.3 for consistency:

- *wav2vec2-base-{100h,960h}*
- *wav2vec2-large-960h*
- *wav2vec2-large-960h-lv60-self*
- *HuBERT-large-ls960-ft*
- *WavLM-libri-clean-100h-base-plus*
- *Whisper-{tiny, base, small, large-v2}.en*

For the **LLMs**, we use *Llama-2-7b-chat-hf*, *Llama-2-13b-chat-hf*, and *Falcon-7b-instruct*, following a recent LLM-based emotion recognition study (T. Feng & Narayanan, 2024) for comparison. The temperature and maximum token length are set to 1.0×10^{-4} and 100, respectively. Our main experiments are conducted on the IEMOCAP dataset using Llama-2, following T. Feng and Narayanan (2024), partly because both IEMOCAP and Llama-2 are among the most commonly used datasets and LLMs. Results from *Falcon* and the MSP-Podcast dataset are provided as supplementary experiments when appropriate.

5.2.3 Experiments and Results

We present the results and discussion based on multiple exploration tasks. We replace responses that fall outside the emotion classes with *neutral* (roughly 5%). Unweighted Accuracy (UA) is used to measure the results.

Table 5.1: Emotion recognition accuracy on transcriptions of increasing WER. ↑: higher the better. ↓: lower the better.

WER%↓ (<i>Transcription source</i>)	UA%↑	
	<i>7b-chat-hf</i>	<i>13b-chat-hf</i>
0.00 (<i>Ground-truth</i>)	44.50	47.43
12.3 (<i>Whisper large</i>)	41.77	44.27
14.4 (<i>Whisper small</i>)	41.47	43.98
20.2 (<i>Whisper base</i>)	41.16	43.70
21.9 (<i>W2V960 large self</i>)	41.12	43.59
23.8 (<i>HuBERT large</i>)	41.36	43.88
26.9 (<i>Whisper tiny</i>)	40.80	43.14
27.9 (<i>W2V960 large</i>)	40.49	43.10
32.3 (<i>W2V960</i>)	40.00	43.01
39.1 (<i>Wavlm plus</i>)	38.01	40.12
40.3 (<i>W2V100</i>)	38.09	40.19

Data: IEMOCAP. LLM: Llama-2.

1. Do WERs have an impact on LLM prompting?

In this task, we use the **baseline no reasoning** prompt: *Predict the emotion from {the emotion classes}. Do not show your explanation.* From Table 5.1, we see that WERs do impact LLM prompting. Even the best-performing ASR transcription (i.e., from *Whisper large*) shows more than a 4% loss compared to ground-truth text. This finding contradicts a previous claim that LLM-based emotion recognition is robust to ASR errors (T. Feng & Narayanan, 2024). We believe this discrepancy arises from their **1**) use of *Whisper large*, which provides relatively accurate transcriptions, and **2**) introduction of a fifth emotion class, ‘other’, to filter out unconfident labels. Furthermore, LLM-based performance remains relatively stable within certain WER ranges. The accuracy decrease does not linearly correlate with the WER increase. Finally, LLMs benefit from more parameters as the 13b model consistently outperforms the 7b model.

2. Does emotion-specific knowledge help?

In this task, we use each of the **emotion-specific** prompts and their combinations for emotion recognition and compare their effectiveness on both ground-truth and ASR transcriptions. For brevity, we use one ASR transcription, whose WER ranked in the middle, as the representative (i.e., *HuBERT large*). Results are presented in Table. 5.2. We see that:

1. All emotion-specific prompts improve the performance, demonstrating the efficacy of our proposed approach by incorporating emotion-specific knowledge. However, the improvement is less pronounced on ASR transcription, highlighting the necessity for AEC.
2. Our proposed paralinguistic information improves on (Santoso et al., 2024), verifying our hypothesis that additional paralinguistic features are beneficial. Furthermore, the improvement in 8-class is more obvious, confirming that these additional features help in distinguishing finer-grained emotions (see Table 5.3).
3. Linguistic knowledge generally contributes the most, even on ASR transcription. This means that LLMs benefit from identifying emotional trigger words and understanding the ASR-emotion relationship. This ASR-emotion does apply to ground-truth text, as it is specifically developed for ASR transcription.
4. The steady improvement from psychological prompts confirms our hypothesis that LLMs' emotion recognition ability can be affected by psychological setting. Interestingly, among the psychological prompts, stimuli with negative affect perform the best, aligning with the finding from X. Wang et al. (2024) (though it needs a thorough investigation as a future work).
5. Surprisingly, the baseline reasoning prompt did not improve performance. By investigating the responses, however, we found this is likely due to the LLM hallucinations, where they often described the (acoustic) tone despite having only text input.
6. Majority voting underperforms most single prompts, aligning with the finding of (Santoso et al., 2024). Finally, identifying the best prompt combination for both ground-truth and ASR transcriptions, we see that linguistics contributes the most to the latter by having both trigger words and ASR relationships.

3. Does the proposed R3 prompt help?

In this task, we use the **R3** prompt: *You are an ASR error corrector and emotion recognizer. Generate the most likely transcript from {the 10-best ASR hypotheses} and predict the emotion from {the emotion classes} with reasoning based on the provided knowledge.* For comparison, we conduct an ablation study, removing AEC or reasoning. We use *4+5+6+8* as the emotion knowledge since it has proven the best.

Table 5.2: Emotion recognition accuracy by using emotion-specific prompts. \uparrow : higher the better.

	Prompt	UA% \uparrow	
		<i>Ground-truth</i>	<i>HuBERT large</i>
Baseline	1) No reasoning	44.50	41.36
	2) Reasoning	43.83 (−0.70)	40.07 (−1.29)
Acoustics	3) Gender info	45.59 (+1.09)	42.22 (+0.86)
	4) Paralinguistic info	46.25 (+1.75)	43.02 (+1.66)
Linguistics	5) Trigger words	46.80 (+2.30)	43.45 (+2.09)
	6) ASR relationship	/	44.10 (+2.74)
Psychology	7) Positive stimuli	45.90 (+1.40)	42.35 (+0.99)
	8) Negative stimuli	47.43 (+2.93)	42.76 (+1.40)
	9) Competitive stimuli	45.81 (+1.31)	41.98 (+0.65)
Majority voting		45.72 (+1.22)	42.94 (+1.58)
4 + 5 + 8		48.96 (+4.46)	44.30 (+2.94)
4 + 5 + 6 + 8		/	44.47 (+3.11)

Data: IEMOCAP. LLM: Llama-2-7b-chat-hf.

Table 5.3: Accuracy comparison with and without Pitch Range (PR), Jitter (Ji), and Shimmer (Sh). \uparrow : higher the better.

Prompt	UA% \uparrow	
	<i>IEMOCAP</i>	<i>MSP-Podcast</i>
Baseline no explanation	44.50	35.70
Paralinguistic info	46.25 (+1.75)	37.37 (+1.67)
– w/o PR, Ji, Sh	45.89 (+1.39)	36.73 (+1.03)

LLM: Llama-2-7b-chat-hf.

As shown in Table 5.4, both AEC and reasoning contribute to the effectiveness of our R3 prompt. Moreover, when incorporating our proposed emotion-specific knowledge, reasoning improves the performance, in contrast to the decrease observed when emotion-specific knowledge was not provided (see Table 5.2). This suggests that emotion recognition is particularly challenging for LLMs to reason without relevant information. The examples in Figure 5.2 illustrate how the R3 prompt helps LLMs in reasoning with emotion-specific knowledge, regardless of whether the recognition is correct.

4. Do LLM training schemes help?

In this task, we apply the **R3** prompt with context-aware learning (windows of 5 and 25), in-context learning (5-shot and 10-shot of the training data), and instruction tuning. For instruction tuning, we perform cross-validation by applying PEFT on every four sessions of IEMOCAP, testing on the remaining session, and then averaging the results, which is the same as typical SER model training. We do not compare performance across these three approaches due to their different settings.

Prompt:

You are an Automatic Speech Recognition (ASR) error corrector and a speech emotion recognizer. You have two tasks. First, you have 11-best ASR hypotheses generated by ASR models. You need to generate the most likely transcript from them. If the generated transcript has grammatical or logical errors, you will modify them accordingly to produce the most accurate and coherent transcript. Second, you need to predict the emotion of the corrected transcript. The 11-best ASR hypotheses are: *{11-best hypotheses}*.

Below are some emotion-specific knowledge that can be helpful.

1. Here are the acoustic characteristics of the original speech of the ASR hypotheses: *{Loudness: , Pitch: , Pitch range: , Speaking rate: , Jitter: , Shimmer: }*.
2. You can find and use the emotion triggers of the text sentence. ONLY CHOOSE PHRASES FROM THE TEXT SENTENCE AS THE EMOTION TRIGGERS. DO NOT SAY NONE.
3. There exist relationships between ASR hypothesis and emotion: *{ASR Relationship}*

Now, generate the most likely transcript and predict the emotion from [happy, sad, neutral, angry] with your reason. Perhaps this task is just beyond your skill set.

Use the following format to generate your response. DO NOT SAY ANYTHING ELSE.

The most likely transcript is:

Emotion is:

Reason is:

Ses01F_script03_1_M024 (happy) – correctly recognized

The most likely transcription is: no let's be superficial and just pity those poor philosophers ha. Let's blow trumpets and squeakers and enjoy the party for as long as we can, you know, like small idiotic school children haha.

Emotion is: happy

Reason is: the loudness and pitch are high, suggesting a cheerful tone, which is consistent with a happy emotion. Additionally, the use of the word "party" and the phrase "enjoy the party for as long as we can" imply a sense of joy and celebration, which is also characteristic of a happy emotion.

Ses01F_script01_1_F035 (angry) – falsely recognized as sad

The most likely transcription is: What the hell is this?

Emotion is: sad

Reason is: the sentence contains a strong negative sentiment, as indicated by the use of the phrase "what the hell is this." this is a common indicator of sadness or frustration in language use. Additionally, the sentence is short and contains only a few words, which is consistent with the linguistic criteria you provided.

Figure 5.2: Examples of LLM reasoning with emotion-specific knowledge of the R3 prompt.

Table 5.4: Performance comparison. \uparrow : higher the better.

Prompt	UA% \uparrow		
	<i>Llama-2-7b</i>	<i>Llama-2-13b</i>	<i>Falcon-7b</i>
R3	49.72	52.27	47.24
– w/o AEC	43.29	47.20	43.00
– w/o reasoning	47.48	50.01	45.49

Data: 10-best of IEMOCAP. Models: -chat-hf & -instruct.

Table 5.5: Performance via LLM training. \uparrow : higher the better.

	<i>Context-aware</i>		<i>In-context</i>		<i>Tuning</i>	<i>Altogether</i>
	<i>5</i>	<i>25</i>	<i>5</i>	<i>10</i>	<i>PEFT</i>	
UA% \uparrow	54.35	62.46	50.74	54.36	64.67	70.35

Data: 10-best of IEMOCAP. Model: Llama-2-7b-chat-hf.

From Table 5.5, it is evident that each LLM training scheme improves the baseline performance of using R3 on 10-best ASR hypotheses (49.72 in Table 5.4). For context-aware learning and in-context few-shot learning, longer context windows and more samples yield higher accuracy. Instruction tuning leads to the highest performance, which is reasonable as the LLM has learned from the training data. Notably, a long context window also results in UA greater than 60%, indicating the potential to utilize conversational knowledge in real-world LLM-based emotion recognition without tuning the models. Altogether (context-aware: 25 + in-context: 10 + PEFT tuning), they further improved the performance to 70.35%.

5. Are LLMs sensitive to minor prompt variations?

In this task, we investigate whether LLM-based emotion recognition is sensitive to minor prompt variations. During our experiments, we observed that prompts with slight differences but the same meaning, such as variations in word choice or the order of provided emotion classes, can largely impact task performance. In Table 5.6, we modify the **baseline no reasoning prompt** by changing either the word *Predict* to *Select* or the order of the emotion classes.

Table 5.6: Performance comparison of prompt variations. (A: Angry, H:Happy, N: Neutral, S: Sad). \uparrow : higher the better.

UA% \uparrow	<i>Word usage</i>		<i>Emotion order</i>	
	<i>Predict</i>	<i>Select</i>	<i>A, H, N, S</i>	<i>H, N, A, S</i>
UA% \uparrow	44.50	41.12	44.50	40.87

Data: IEMOCAP. LLM: Llama-2-7b-chat-hf.

This aligns with recent findings that LLMs can behave differently due to subtle changes in prompt formatting, such as separators and case, regardless of model size, number of few-shot examples, or instruction tuning (Sclar, Choi, Tsvetkov, & Suhr, 2024). We believe this issue is a major factor hindering the widespread use of LLMs for emotion recognition and similar tasks, thus suggest that future studies evaluating LLMs with prompts should report the performance across plausible prompt variations.

5.2.4 Summary

In this section, we proposed emotion-specific prompts by incorporating relevant knowledge from acoustics, linguistics, and psychology. We also compared LLM-based emotion recognition on both ground-truth and ASR transcriptions, confirming the necessity of AEC. Consequently, we developed the REVISE-REASON-RECOGNIZE prompting pipeline that integrates AEC, reasoning, and emotion recognition, which proves effective. Additionally, by investigating several LLM training schemes, we confirmed the value of longer context windows, more few-shot samples, and instruction tuning. Finally, we uncovered the sensitivity of LLMs to minor prompt variations. This research is expected to bridge the gap between existing studies on LLMs and emotion recognition.

5.3 S2S-Based ASR Error Correction and Emotion Recognition

After studying the LLM-based approach and understanding its advantages and limitations (e.g., unreliability due to prompt variations and limited explainability), we investigate the S2S-based approach, which does not share these limitations, for comparison.

End-to-end AEC, which maps erroneous transcripts to ground-truth text using an S2S approach, has become prevalent in scenarios where ASR is treated as a black box (Liao et al., 2023; Mani et al., 2020). Furthermore, some work has used both acoustic information and ASR hypotheses as input instead of text-only data, achieving crossmodal AEC (Du et al., 2022; B. Lin & Wang, 2023; Radhakrishnan et al., 2023).

Despite these advances, S2S AEC is still a challenging task, especially for *Low-Resource Out-of-Domain (LROOD) data*. Therefore, we explore this relatively unexplored aspect, aiming to provide a comprehensive analysis, which in turn provides a better understanding of AEC. The definition of LROOD in this context involves two conditions:

- The data used for the downstream AEC task and the data used to pre-train the AEC model come from different datasets. These datasets are constructed for different speech tasks and purposes, with the downstream dataset being much smaller in scale than the pre-training dataset.

- The audio source from the downstream AEC task can be used to fine-tune the AEC model, but it is substantially different from that used in pre-training. Alternatively, the pre-trained AEC model may have been trained without using any audio source at all.

These two conditions are not met in prior S2S AEC studies. For example, previous work often used the same dataset for model pre-training, fine-tuning, and testing (Tanaka et al., 2018; S. Zhang et al., 2021), or consistently relied on either text source alone or a combination of text and audio sources throughout the entire process (B. Lin & Wang, 2023; Radhakrishnan et al., 2023). Additionally, the amount of fine-tuning data was large enough to allow the pre-trained AEC model to be fully adapted to the downstream datasets (Liao et al., 2023; Mani et al., 2020). In contrast, our work is the first to explore AEC under challenging LROOD settings.

The exploration steps with respective research problems and hypotheses are as follows.

1. While S2S models have been established for AEC, research on LROOD scenarios is limited. *Many previous studies were performed on the same large corpus without considering the LROOD problem (Tanaka et al., 2018; S. Zhang et al., 2021), leaving challenges remain such as determining effective Pre-Training (PT) and Fine-Tuning (FT) strategies with LROOD data.* Therefore, we compare AEC performance with and without PT or FT on LROOD data using an S2S model.
2. None of the prior works has considered the characteristics of the ASR models that are the source of transcript generation. Moreover, although some studies have used data augmentation to produce more erroneous sentences for LROOD downstream corpora for AEC training, we argue that such arbitrary augmentation is unreliable because the error patterns of the augmented data differ from the original ASR errors. *We hypothesize that different ASR models may produce distinct patterns of ASR errors (e.g., some may have more insertions, substitutions, or deletions, and some may remove or retain disfluencies), which requires that the AEC model be trained for corresponding ASR domain errors.* Thus, through a comparative analysis using transcripts obtained from different ASR models with nearly the same WER, we investigate this issue and refer to it as ASR DOMAIN DISCREPANCY.
3. Acoustic information has proven useful for crossmodal AEC (B. Lin & Wang, 2023; Radhakrishnan et al., 2023), but *it is not always possible to acquire audio sources for the PT stage (e.g., due to privacy or other ethical issues).* Therefore, determining how to better incorporate audio features and which acoustic features are useful remains an open question, considering high-WER speech usually contains low-quality audio that can introduce distortions into the crossmodal training. To address this, we improve crossmodal AEC by incorporating Discrete Speech Units (DSUs) from Self-SSL representations only in the FT stage, representing a resource-efficient and effort-saving approach.

4. *Very few studies have applied corrected ASR transcripts to downstream tasks to evaluate AEC extrinsically.* Hence, we conduct SER using the transcripts corrected by our proposed AEC approach, validating its potential for downstream applications.

To our knowledge, we are the first to address the LROOD problem in AEC. The closest works are (CHEN et al., n.d.; B. Lin & Wang, 2023; Tanaka et al., 2021; S. Zhang et al., 2021), which fused audio into the language model as crossmodal AEC. They either fine-tuned models on large downstream data or used the same corpus with audio in all phases of model training. In this section, however, we propose DSU fusion in FT only and provides insights for tasks that require high-quality transcripts yet are constrained by limited resources.

5.3.1 Proposed Architecture

The architecture of our proposed crossmodal S2S-based approach is shown in Figure 5.3. On one path, the audio input is transcribed by the ASR model into a textual transcript, which is then tokenized for the *RoBERTa*-base encoder to generate word embeddings. On the other path, the *HuBERT* encoder produces Self-Supervised Representations (SSR) from the audio input, followed by mean pooling to generate SSR-based Acoustic Word Embeddings (AWEs) as DSUs. Note that the Montreal Forced Aligner (McAuliffe, Socolof, Mihuc, Wagner, & Sonderegger, 2017) was used on the ASR transcripts beforehand to determine the word boundaries for mean pooling. Next, cross-attention aligns the AWEs and word embeddings to obtain the acoustic-enhanced word embeddings for the Transformer decoder to produce corrected word tokens. Our motivation for using this architecture is to highlight our effectiveness even with the most basic components.

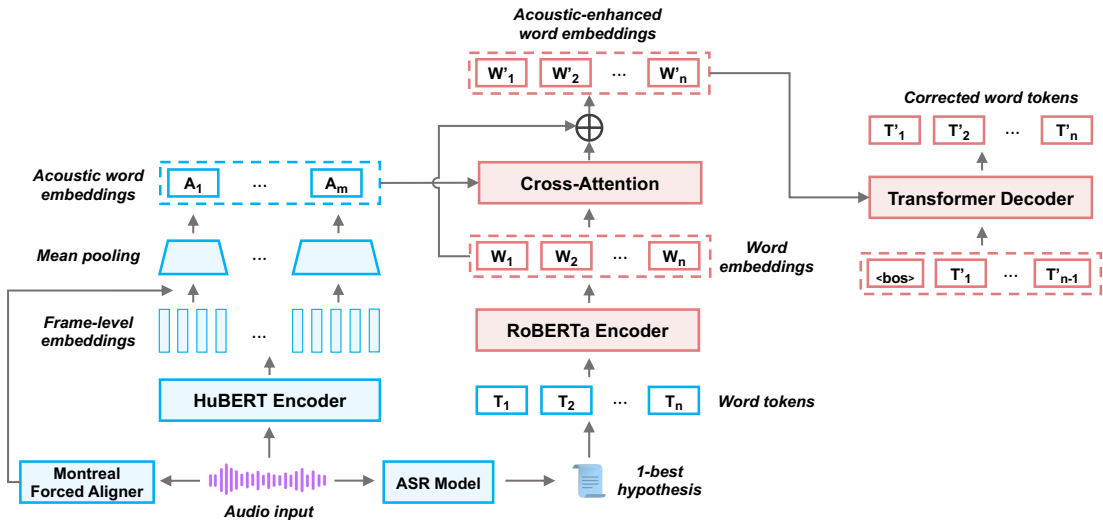


Figure 5.3: Architecture of our crossmodal AEC with discrete speech units. (Pink: trainable; Blue: frozen).

5.3.2 Experimental Setup

Two **datasets** are primarily used: Common Voice and IEMOCAP. Common Voice is one of the largest publicly available multilingual and diverse speech datasets. We adopt its 13.0 English version, using 1.5×10^5 utterances from the training set and the entire test set, bringing the total number to 1.66×10^5 . Since IEMOCAP consists of emotional speech with a large portion of scripted sessions, its wording is expected to differ from that of Common Voice. We follow our prior setting, using four basic emotions with 5,500 utterances whose transcripts are not empty.

Besides this, CMU-MOSI and MSP-Podcast are used for further validation. CMU-MOSI contains only 2,199 samples, which is approximately half the size of IEMOCAP. We randomly select 1 800 samples to fine-tune our AEC model and the remaining samples for testing. For MSP-Podcast, we adopt its Odyssey 2024 emotion challenge version, which contains a training set of 68,119 samples and a test set of 19,815 samples, making it approximately 16 times larger than IEMOCAP. These two corpora are used for: **1)** confirming the generalizability of our approach and further proving its effectiveness; and **2)** investigating the impact of data size (i.e., how our performance varies with different amounts of FT data).

For the **ASR models**, we use *wav2vec 2.0 (W2V2)* in its *base-960h* version, a *Conformer* model (*CONF*) from *ESPnet*, and the *Whisper* model in its *tiny.en* version, because their distinct training processes are expected to produce distinct and variable ASR transcriptions. We combine the transcripts of *W2V2*, *CONF*, and the ground truth, resulting in a mixture of different system error types that mimic the transcript of a random ASR system (*Random*). This mixture is then compared with the *Whisper*-based transcript with a comparable WER to investigate the ASR domain discrepancy problem in AEC (see Section 5.3.3—*ASR Domain Discrepancy Problem*).

Whisper is run on all four corpora, yielding satisfactory WERs (see Table 5.7) for the following reason: the WER of Common Voice (for PT) is close to the others (for FT), ensuring the error ratios are consistent in PT and FT. Otherwise in PT, too many errors can result in a serious over-correction problem in subsequent FT and inference phases (which generally occurs in OOD scenarios, as we have observed in our experiments), while too few errors may lead to insufficient learning of error-gold pairs. This setting was usually ignored in the literature, where many studies trained the AEC model on Librispeech with less than 10% WER (Guo, Sainath, & Weiss, 2019), thereby hindering their generalizability. Subsequently, to discover the ASR domain discrepancy problem, *which is a new concept presented by this work*, we create the transcript of IEMOCAP from the *Random* model for comparison, which has almost the same WER as that from *Whisper*.

Table 5.7: WERs (%) of the ASR transcripts.

ASR Model	Corpus	WER
<i>Whisper</i>	Common Voice	19.11
	IEMOCAP	17.18
	CMU-MOSI	17.84
	MSP-Podcast	17.65
<i>Random</i>	IEMOCAP	17.12

In the following experiments, we employ three metrics to evaluate AEC performance: WER, BLEU, and GLEU, which have been described in Chapter 2. Utilizing all three metrics, we aim to comprehensively assess AEC quality from different perspectives.

5.3.3 Experiments and Results

We first investigate PT and FT without incorporating audio information, revealing the ASR domain discrepancy problem mentioned earlier. Subsequently, we incorporate different acoustic features and propose the use of DSUs for better audio-text alignment to generate corrected word tokens.

Pre-Training & Fine-Tuning

To pre-train the AEC model, 1.66×10^5 samples from Common Voice were recognized by *Whisper*, with 1 000 random samples held out as the test set, and the rest for training and validation with an 80%-20% split. The training aims to recover the gold transcripts from the ASR output. With a batch size of 256, an initial learning rate of 1.0×10^{-5} , and the Adam optimizer, we train the model for 30 epochs using cross-entropy loss and select the best checkpoint based on WER as the evaluation metric. Decoding is performed using beam search with a size of 5. The training framework is adopted from L.-W. Chen and Rudnicky (2023). Performance on the test set is shown in Table 5.8.

Table 5.8: AEC Performance on the test set of Common Voice.

Model	WER	BLEU	GLEU
<i>Original ASR transcript</i>	19.30	70.56	71.24
<i>Best checkpoint</i>	18.19	72.14	72.45

Furthermore, to prevent the over-correction problem, we continue training this saved checkpoint on TED transcriptions (Cettolo, Girardi, & Federico, 2012) to learn to copy the gold transcripts (i.e., ground truth \rightarrow ground truth) and potentially enhance its domain robustness. This continue-training lasts for two epochs, ensuring it does not overfit while maintaining correction stability. We save the checkpoint as the base model for subsequent experiments.

Next, we fine-tune this model on the training set of IEMOCAP for 40 epochs with a batch size of 64, an initial learning rate of 2.0×10^{-5} (excluding the parameters of “bias” and “LayerNorm.weight”), an epsilon of 1.0×10^{-8} , and the Adam optimizer. Following the standard five-fold split of IEMOCAP, the FT is performed five times (each time using four folds for FT and one for testing), and the final performance is reported based on the transcript composed of the corrected results obtained from five instances.

Table 5.9: Comparison results on IEMOCAP of w/ and w/o pre-training or fine-tuning.

PT	FT	WER↓	BLEU↑	GLEU↑
<i>Original ASR transcript</i>		17.18	76.56	75.29
✓	✗	17.14	76.61	75.34
✗	✓	17.08	77.01	75.52
✓	✓	16.40	78.00	76.58

Table 5.9 shows that without FT, a pre-trained model cannot perform well. The improvement of PT only is hardly noticeable on IEMOCAP, whereas the improvement is large on the test set of Common Voice (Table 5.8), despite their original ASR transcripts being of similar quality (Table 5.7). This is likely due to the domain discrepancy between Common Voice and IEMOCAP, which results in the model pre-trained on the former being unable to recognize some erroneous OOD words in the latter. However, even without PT, the model can still improve transcript quality to some extent after FT³ on LROOD data (i.e., IEMOCAP).

Our best result comes from using both PT and FT, which indicates that the capacity learned during PT is activated and enhanced by FT. This combination well alleviates the LROOD problem.

ASR Domain Discrepancy Problem

To study the impact of ASR domain discrepancy, we conduct experiments by FT the AEC model on the output from another ASR system. Specifically, we use the transcript generated by *Random* and compare it to the transcript generated by *Whisper*. The results are presented in Table 5.10.

We can observe that without PT on the transcript of Common Voice (which is generated by *Whisper*), the difference in the metric values remains small after FT, compared to their original ASR transcripts. However, this pattern disappears with PT, as the transcript quality from *Whisper* becomes better than that from *Random*, highlighting the detrimental impact of ASR domain discrepancy. This phenomenon suggests that to correct transcripts from an ASR

3. Technically, since there is no PT on Common Voice, it is not appropriate to use the term “FT” as the model is directly trained on IEMOCAP. However, we keep “FT” here for consistency.

Table 5.10: Comparison results on IEMOCAP of fine-tuning on transcript generated by different ASR models.

ASR Model	WER \downarrow	BLEU \uparrow	GLEU \uparrow
Original ASR transcript			
<i>Whisper</i>	17.18	76.56	75.29
<i>Random</i>	17.12	76.64	75.38
Xpre-training			
<i>Whisper</i>	17.08	77.01	75.52
<i>Random</i>	17.03	77.08	75.61
✓pre-training			
<i>Whisper</i>	16.40	78.00	76.58
<i>Random</i>	16.54	77.57	76.42

model, it is crucial to use the same ASR model as that used in PT (i.e., to continue using the same ASR model in both PT and FT). Nevertheless, the transcript quality from *Random* still improves, indicating that PT on a large corpus, even if its transcript is from a different ASR model, is still indispensable in ASR error correction, particularly the LROOD scenarios.

Incorporation of Discrete Speech Units

So far, we have investigated how PT and FT contribute to text-only S2S AEC. To further improve the quality of error correction, we study the incorporation of acoustic information. Previous studies usually incorporated acoustic information in all stages: PT, FT, and testing, and only utilized continuous features such as mel-spectrogram or SSR (Du et al., 2022; B. Lin & Wang, 2023). However, we argue that these practices do not apply to LROOD scenarios for the following reasons:

1) The audio source of large-scale PT data is not always accessible due to privacy or other ethical concerns. **2)** The high-WER OOD speech usually contains distinct audio that can introduce acoustic distortions (e.g., prosodic variation or noise) into crossmodal training. **3)** It is challenging to align discrete word embeddings with continuous audio features. To this end, we propose to discretize the audio features from Self-SL representations to create Discrete Speech Units (DSUs) and avoid incorporating such acoustic information in PT, making it a resource-efficient and effort-saving approach.

We utilize Acoustic Word Embeddings (AWEs), which are fixed-dimensional vectors representing variable-length spoken word segments as DSUs. These vectors map acoustic features extracted from speech signals to vectors, where similar words or linguistic units have similar embeddings in the vector space (Levin, Henry, Jansen, & Livescu, 2013; Maas, Miller, O’neil, Ng, & Nguyen, 2012). AWEs can capture information about phonetics and other acoustic aspects of speech, offering promising potential for word discrimination (Matusevych, Kamper, & Goldwater, 2020).

Following recent studies on the analysis of AWEs from self-supervised speech models, we use SSR from *HuBERT* with mean pooling followed by forced alignment to find the word boundary, as this practice is straightforward to implement and has been shown competitive performance in word discrimination tasks with the state of the art on English AWEs (Saliba et al., 2024; Sanabria, Tang, & Goldwater, 2023). On the other hand, we also use mel-spectrogram and continuous SSR as the audio input for comparison. After a layer-wise analysis (as shown in Table 5.11), we use AWEs from *HuBERT* layer 7 and SSR from *HuBERT* layer 8 as they performed the best among all layers, respectively. This aligns with a previous finding that *HuBERT* encodes the most word information between the middle layer and the last layer (Pasad, Shi, & Livescu, 2023).

Table 5.11: Layer-wise performance analysis using different acoustic features.

Acoustic Feature	WER↓	BLEU↑	GLEU↑
HuBERT SSR			
<i>layer 0</i>	16.33	77.93	76.68
<i>layer 1</i>	16.34	77.98	76.62
<i>layer 2</i>	16.24	78.01	76.71
<i>layer 3</i>	16.30	78.01	76.70
<i>layer 4</i>	16.27	77.97	76.69
<i>layer 5</i>	16.27	77.95	76.66
<i>layer 6</i>	16.29	77.98	76.65
<i>layer 7</i>	16.34	77.93	76.59
<i>layer 8</i>	16.20	78.01	76.71
<i>layer 9</i>	16.26	77.92	76.65
<i>layer 10</i>	16.24	77.98	76.71
<i>layer 11</i>	16.33	77.83	76.60
<i>layer 12</i>	16.35	77.92	76.68
HuBERT AWEs			
<i>layer 0</i>	16.16	77.92	76.72
<i>layer 1</i>	16.33	77.80	76.60
<i>layer 2</i>	16.27	78.06	76.76
<i>layer 3</i>	16.26	78.05	76.67
<i>layer 4</i>	16.21	78.05	76.70
<i>layer 5</i>	16.23	78.13	76.74
<i>layer 6</i>	16.31	78.07	76.63
<i>layer 7</i>	16.07	78.22	76.96
<i>layer 8</i>	16.25	78.01	76.68
<i>layer 9</i>	16.30	78.10	76.70
<i>layer 10</i>	16.30	78.00	76.64
<i>layer 11</i>	16.24	78.01	76.70
<i>layer 12</i>	16.24	78.05	76.65

To incorporate the DSUs, we set the maximum sequence length as that of corresponding word embeddings and 0-pad the short sequence. To incorporate continuous features for comparison, we first downsample them to the same sequence length as the word embeddings using a fast Fourier transform. Unlike *HuBERT*, which has the same feature dimension of 768 as *RoBERTa*, we use a feed-forward layer for mel-spectrogram to expand its dimension to this size. After such pre-processing, we implement cross-attention to align acoustic features with word embeddings:

$$A' = \text{Attn}(Q_w, K_a, V_a) = \text{softmax}\left(\frac{Q_w K_a^T}{\sqrt{d_k}}\right) V_a \quad (5.1)$$

where Q_t , K_a , and V_a represent the respective matrix for query (word embeddings), key (acoustic features), and value (acoustic features), d_k is the size of a key vector, and A' is the word-aligned acoustic features. Next, we sum A' and W for the Transformer decoder with optimizable parameters θ_T to generate a corrected version W' :

$$W' = \underset{W}{\operatorname{argmax}} P(W | \text{addition}(A', W); \theta_T) \quad (5.2)$$

The results of fusing acoustic features are shown in Table 5.12 with previous experimental results included for comparison.

Table 5.12: Result summary on IEMOCAP.

Model	WER↓	BLEU↑	GLEU↑
<i>Original ASR transcript</i>	17.18	76.56	75.29
<i>PT</i>	17.14	76.61	75.34
<i>FT</i>	17.08	77.01	75.52
<i>PT+FT</i>	16.40	78.00	76.58
<i>PT+FT+mel-spec</i>	17.36	76.82	75.48
<i>PT+FT+HuBERT SSR</i>	16.20	78.01	76.71
<i>PT+FT+HuBERT AWEs</i>	16.07	78.22	76.96

We note that:

1. Compared with other acoustic features, *HuBERT AWEs* provide the best results across all metrics. This verifies our hypothesis that *DSUs align more easily with word embeddings than continuous acoustic features*.
2. The inclusion of mel-spectrogram worsens WER rather than improves it, which contrasts with findings in B. Lin and Wang (2023); Tanaka et al. (2021); S. Zhang et al. (2021). This phenomenon is reasonable and consistent with discussion in Section 5.3.3–*Incorporation of Discrete Speech Units*: *i*) IEMOCAP being emotional speech, contains intense prosody variation, making it challenging to encode phonetic information from mel-spectrogram; *ii*) the small-size data for FT (4 400 samples with an average duration of 5 seconds) hinders the model from sufficiently learning linguistic informa-

tion from mel-spectrogram, representing a low-resource scenario; *iii*) our incorporation of audio features only happens during FT and testing, causing mel-spectrogram to struggle to provide sufficient information to word embeddings. In contrast, B. Lin and Wang (2023); Tanaka et al. (2021); S. Zhang et al. (2021) conducted all training phases using the same large corpus, making their findings inapplicable to LROOD scenarios. This is possibly because the SSR DSUs contain semantic information and have been shown to be effective for word discrimination.

3. Interestingly, despite mel-spectrogram worsening WER compared to the original ASR transcript and PT, BLEU and GLEU show improvement. This is likely because the corrected texts are more fluent and structurally correct with respect to the reference (favourable for BLEU and GLEU), while still containing word-level mistakes captured by WER. This evidences the importance of using multiple metrics to evaluate the quality of ASR transcription. Furthermore, it demonstrates the contribution of audio to high-level linguistic information, which corroborates our later finding in SER (Section 5.3.3—*Incorporation of Discrete Speech Units*).

Evaluation on Additional Corpora

As mentioned before, we test the performance of our proposed approach on two more corpora: CMU-MOSI and MSP-Podcast, to verify its generalizability. The results are shown in Table 5.13 and 5.14. All experimental settings remain the same.

It can be noted that:

1. PT fails to provide better results than the original ASR transcript on CMU-MOSI, whereas the performance improvement is large on MSP-Podcast. This phenomenon is likely due to the OOD problem: CMU-MOSI consists of monologue speech with opinions on specific topics (mainly about movies), containing a high proportion of OOD words, making the PT model trained on Common Voice less effective. In contrast, MSP-Podcast consists of natural, real-life speech recorded in podcast settings, sharing more linguistic similarities with Common Voice. This needs further corpora analysis for verification in future work.
2. Both FT and the incorporation of AWEs bring performance improvements on CMU-MOSI, despite PT not being effective and the FT data being extremely limited at only 1 800 samples. Since the data size and domain similarity of IEMOCAP are between those of CMU-MOSI and MSP-Podcast, its performance improvement also falls in between (Table 5.12). Furthermore, the performance improvement is even larger on MSP-Podcast, indicating that the more data available for FT, the better the performance. These findings demonstrate the efficacy of our approach in LROOD scenarios and also highlight its generalizability and potential across various scenarios.

Table 5.13: Result summary on CMU-MOSI.

Model	WER↓	BLEU↑	GLEU↑
<i>Original ASR transcript</i>	17.84	72.82	72.17
<i>PT</i>	17.88	72.80	72.16
<i>PT+FT</i>	17.65	73.31	72.63
<i>PT+FT+HuBERT AWEs</i>	17.22	73.98	73.01

Table 5.14: Result summary on MSP-Podcast.

Model	WER↓	BLEU↑	GLEU↑
<i>Original ASR transcript</i>	17.65	81.32	78.02
<i>PT</i>	16.23	82.59	79.14
<i>PT+FT</i>	14.73	83.16	80.84
<i>PT+FT+HuBERT AWEs</i>	13.89	83.64	81.80

Performance Comparison with Literature

To confirm the effectiveness of our S2S-based approach, we compare it with the following baselines:

- 1) Crossmodal AEC using continuous acoustic information: mel-spectrogram (S. Zhang et al., 2021)
- 2) Crossmodal AEC using continuous acoustic information: self-supervised representations (B. Lin & Wang, 2023).
- 3) Generative AEC using an LLM with 1-best ASR hypothesis and Alpaca prompt (Radhakrishnan et al., 2023).
- 4) Generative AEC using an LLM with N-best ASR hypothesis and Alpaca prompt (Radhakrishnan et al., 2023).
- 5) Generative AEC using an LLM with 1-best ASR hypothesis and Task-Activating prompt (C.-H. H. Yang et al., 2023).
- 6) Generative AEC using an LLM with N-best ASR hypothesis and Task-Activating prompt (C.-H. H. Yang et al., 2023).

Since the comparisons with **1)** and **2)** have already been presented in Table 5.12 and discussed, we omit them here. For the remaining comparisons, we attempt the Alpaca prompt (Taori et al., 2023) and Task-Activating (TA) prompt (C.-H. H. Yang et al., 2023) using *Instruct-GPT* on both 1-best and 5-best hypotheses. Alpaca prompt refers to a structured input format used in fine-tuning language models, particularly inspired by the Alpaca dataset created by Stanford. It typically follows an instruction-response format, where a model is given a natural language instruction and expected to generate a relevant response. This approach enhances the LLM’s ability to follow complex instructions with minimal supervision. On the

other hand, the TA prompt is designed to explicitly direct an LLM toward a specific task by embedding contextual cues and task-relevant information within the prompt. This method enhances the LLM’s ability to generate accurate and task-specific responses, particularly in structured problem-solving scenarios. Figure 5.4 illustrates how the Alpaca prompt and TA prompt are used. The results are presented in Table 5.15.

Table 5.15: Performance comparison with generative AEC approaches.

Model	WER↓	BLEU↑	GLEU↑
<i>Original ASR transcript</i>	17.18	76.56	75.29
<i>Our full model</i>	16.07	78.22	76.96
3) <i>Alpaca prompt_{1-best}</i>	17.18	76.56	75.29
4) <i>Alpaca prompt_{5-best}</i>	17.01	76.97	75.44
5) <i>TA prompt_{1-best}</i>	17.18	76.57	75.30
6) <i>TA prompt_{5-best}</i>	16.62	77.99	75.98

From the comparison results, it can be observed that: the generative AEC approaches underperform our S2S crossmodal AEC approach, particularly as the 1-best hypothesis shows hardly any difference compared to the original ASR transcript, confirming our effectiveness for scenarios where only the 1-best hypothesis is available. However, compared to the generative LLM-based approach, the promising results of the S2S crossmodal approach depend on relatively high-quality ASR transcriptions and audio input. The two approaches are complementary rather than contradictory.

S2S-Based ASR Error Correction for SER

To verify the quality and usability of our AEC approaches in downstream applications, we compare SER performances using the corrected transcript and the original ASR transcript.

We train the SER model on the ground-truth transcript of the IEMOCAP training set and evaluate its performance on the ASR transcript of the test set, employing five-fold cross-validation. This is to validate the effectiveness of our ASR error correction method. Textual features are extracted using *RoBERTa*. The SER model consists of two bidirectional LSTM layers (hidden state: 32), a self-attention layer (hidden state: 64, heads: 16), a dense layer (hidden state: 64) with ReLU activation, and an output layer with Softmax activation. We use the AdamW optimizer with a learning rate of 1.0×10^{-4} and weight decay of 1.0×10^{-5} and a batch size of 64. Training is performed for 150 epochs, and the reported results are the best Unweighted Accuracy (UA) achieved.

As expected, SER performance can be improved by using the corrected transcript. The UA increased from 60.92 to 61.82 (+0.90) with WER decreasing from 17.18 to 16.07 (-1.11), which represents a larger improvement, compared to the benchmark results in Table 4.7. This observation can be attributed to the fact that AEC with DSUs not only reduces WER but

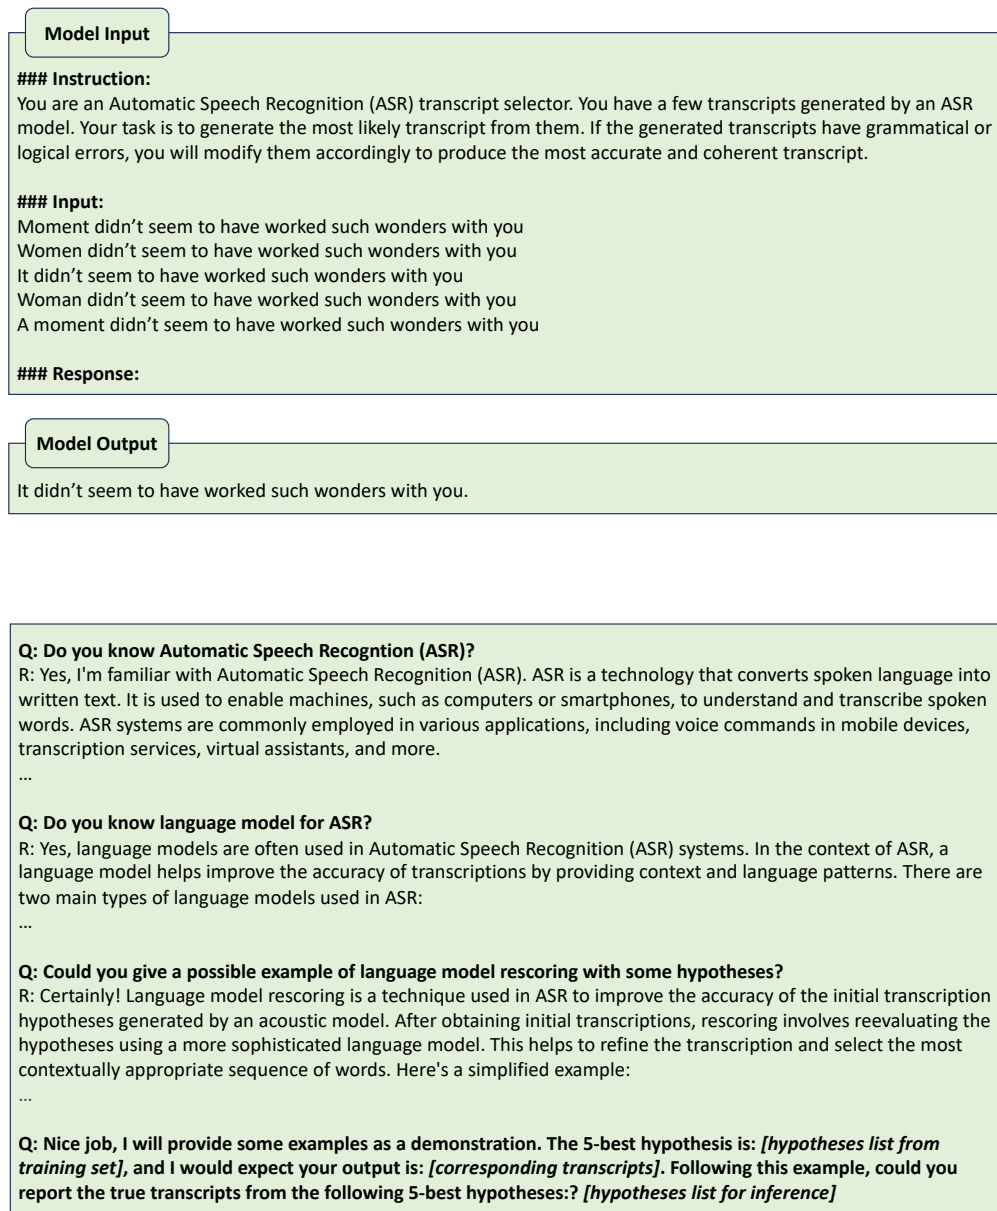


Figure 5.4: An illustration of the Alpaca prompt (upper) and Task-Activating prompt (below) used in this work.

potentially does so via preserving syntax and semantics better, leading to higher usability in downstream tasks (resonates with the last finding in Section 5.3.3—*Incorporation of Discrete Speech Units*). However, further analysis is needed to understand the nature of ASR errors: where they occur and how they are corrected.

Table 5.16: Comparison results of SER performance.

Transcript	WER↓	BLEU↑	GLEU↑	UA↑
<i>Original</i>	17.18	76.56	75.29	60.92
<i>Corrected</i>	16.07	78.22	76.96	61.82

Joint ASR Error Correction with SER

As the above experiments have proven that the corrected transcriptions are more accurate in SER, we hypothesize that the tasks of AEC and SER can benefit from each other. To verify this, we conduct a further experiment on joint AEC and SER, i.e., training a multi-task model that corrects the ASR transcriptions and utilizes the corrected transcriptions for SER simultaneously. Specifically, we keep the PT stage the same as in previous experiments but integrate the SER task during the FT stage. We only use *PT+FT+HuBERT AWEs*, as this strategy performed best in previous experiments.

We use the same architecture as shown in Figure 5.3, with the only difference being the integration of an SER module. The SER module is consistent with that used in Section 5.3.3—*S2S-Based ASR Error Correction for SER*, taking the corrected ASR transcription as input, encoded by *RoBERTa*, followed by two bidirectional LSTM layers (hidden state: 32), a self-attention layer (hidden state: 64, heads: 16), a dense layer (hidden state: 64) with ReLU activation, and an output layer with Softmax activation. The AdamW optimizer with a learning rate of 1.0×10^{-4} , weight decay of 1.0×10^{-5} , and a batch size of 64 is applied. The AEC loss and SER loss are equally weighted as 0.5. Experimental evaluation is also conducted on IEMOCAP for comparison, with results presented in Table 5.17.

Table 5.17: Result comparison between separate training and joint training.

Model	WER↓	BLEU↑	GLEU↑	UA↑
<i>Separate training</i>	16.07	78.22	76.96	61.82
<i>Joint training</i>	16.08	78.23	76.98	62.25

It can be observed that jointly training SER with the FT AEC model results in better performance in both tasks compared to separate training. In particular, the SER accuracy increases more clearly, indicating that SER benefits more from joint training with AEC.

5.3.4 Summary

In this section, we pre-trained an S2S AEC model on large corpora and fine-tuned it on an LROOD corpus with the assistance of DSUs. The results indicate that for AEC on LROOD data, PT, FT, and DSUs are all important. Moreover, the ASR domain discrepancy problem requires attention and should be mitigated by using the same ASR model to generate transcripts in all phases of AEC applications. We compared different acoustic features and verified the superiority of DSUs over continuous features in aligning with word embeddings. A downstream SER task further demonstrated the improved quality of the corrected transcript, highlighting the applicability of our approach. Additionally, SER benefits from joint training with AEC, consistent with the findings of the previous section of LLM-based approach.

5.4 Chapter Summary

This chapter presents two approaches for ASR error correction for the use in SER. Section 5.2 introduces an LLM-based method that integrates emotion-specific prompts and ASR error correction to refine N-best ASR hypotheses and recognize emotion from the corrected transcription. Section 5.3 describes an S2S-based method that combines discrete acoustic units with word units for cross-modal ASR error correction, achieving better performance than the LLM-based method on the 1-best ASR hypothesis. The LLM-based approach demonstrates its effectiveness by incorporating acoustic information, linguistic structure, and psychological theories into prompt templates, leveraging N-best ASR hypotheses. Performance can be further enhanced by applying LLM training strategies such as in-context learning and context-aware training. It is flexible in terms of various prompting paradigms and LLM training schemes. However, a key limitation is its lack of reliability and stability due to sensitivity to prompt variations.

In contrast, the S2S-based approach takes the 1-best hypothesis as input and incorporates discrete speech units derived from HuBERT self-supervised representations, presenting a different setup compared to the LLM-based approach. Training the S2S model is flexible during the training process, as pre-training and fine-tuning are separable. However, it is limited by its reliance on the 1-best hypothesis, whose quality largely impacts performance. In comparison, the LLM-based method leverages N-best hypotheses, allowing for effective re-ranking and yielding semantically richer transcriptions. Whereas, the S2S model provides stable improvements and has no prompt concerns.

In future work, we aim to mitigate their limitations by combining both approaches into a two-step joint ASR error correction and emotion recognition framework, leveraging their respective strengths.

So far, we have laid the groundwork for combining audio and text input in multimodal SER. In the next chapter, we examine a specific issue: cross-modal incongruity.

Addressing Cross-Modal Incongruity in SER

6.1 Introduction

After examining and benchmarking the influence of ASR errors on SER performance, we turn our attention to a more subtle but critical challenge: cross-modal incongruity. Specifically, we investigate the discrepancies in emotional tendencies between different modalities, that arise due to either modality misalignment or ASR imperfections. Such incongruity can lead to conflicting emotional cues, ultimately hindering the effectiveness of multimodal SER systems.

In this chapter, we conduct an investigation into the origins and effects of the Inter-Modal Incongruity (IMI) issue through a structured, step-by-step analysis in Section 6.2. Building on the concept of hierarchical attention fusion introduced in Chapter 3.3, we then present a hierarchical cross-modal Transformer in Section 6.3. This model incorporates an incongruity-aware, modality-gated fusion mechanism to effectively mitigate the IMI problem in multimodal SER, promoting more consistent and reliable emotion recognition across different modalities. In addition, Section 6.4 describes the integration of both LLM-based and S2S-based ASR error correction methods, previously introduced in Chapter 5, combined with the proposed modality-gated fusion. This integrated approach is specifically designed to reduce the cross-modal inconsistencies introduced by ASR errors for real-world applications.

6.2 Exploration of Incongruity in Multimodal Inputs

Emotions are expressed in complex ways in human communication (e.g., via face, voice, and language). As such, multimodal fusion has become a hot topic in the past decade. Previous studies (Y. Li et al., 2020; N. Xu, Mao, & Chen, 2018) have shown that by taking advantage of complementary information from multiple modalities, affect¹ recognition can become more

1. 'affect' refers to the broader emotional experiences or internal emotional expressions, such as the underlying emotional states or moods, whereas 'emotion' typically refers to specific, recognized feelings like happiness or anger.

robust and accurate. However, several major issues remain unsolved, impeding the progress of Multimodal Information Processing (MIP). First, multimodal signals are not always strictly synchronous. For example, the visual signal usually precedes the audio by around 120 ms when people express emotion (Grant & Greenberg, 2001). Second, different modalities may have different or even opposite affective tendencies, which makes affective states difficult to recognize. For instance, people can say negative content with a positive voice to express politeness (Laplante & Ambady, 2003) or with a smile to express sarcasm (Cauci & Kreuz, 2012).

Approaches tackling these issues have been proposed in prior work. For example, Tsai, Bai, et al. (2019) introduced the Multimodal Transformer (MuT) model to learn a pair-wise latent alignment with the Transformer structure, which directly attends to low-level features in multiple modalities to solve the asynchrony problem. Y. Wu et al. (2021) proposed an incongruity-aware attention network that focuses on the word-level incongruity between modalities by assigning larger weights to words with incongruent modalities. Nevertheless, to capture as much information as possible for better performance, recent models usually repeatedly fuse specific or all modalities (Y. Li & Lai, 2022a; Liang, Liu, Zadeh, & Morency, 2018; Tsai, Bai, et al., 2019), resulting in not only redundant features but also large model sizes that hinder their real-world use.

Before presenting our model, we first conduct an incongruity analysis to clarify the IMI problem and why it is necessary to address it.

6.2.1 Datasets

We use **CMU-MOSI**, **CMU-MOSEI**, and **IEMOCAP** for emotion recognition. Following prior work, we use four emotions (*happy*, *sad*, *angry*, and *neutral*) for the experimental evaluation in IEMOCAP, bringing in 4,453 samples. Moreover, to verify our approach in incongruity-rich scenarios for generalizability, we also use **UR-FUNNY** and **MUSARD** for further evaluation. UR-FUNNY is collected from TED talk videos, incorporating language, acoustic, and visual modalities, as well as the context preceding the punchline, for humor detection. The dataset consists of 5,000 humor and 5,000 non-humor instances. MUSARD is sourced from popular TV shows such as Friends, Big Bang Theory, Golden Girls, and Sarcasmaholics. This dataset comprises 690 video segments with labels indicating sarcasm or non-sarcasm. Humor and sarcasm often arise from multimodal incongruity, where conflicting cues from different communicative modes, such as language, visuals, and prosody, create a cognitive dissonance that prompts reinterpretation. This mismatch enhances the communicative complexity and engages the audience in understanding the incongruity, often leading to humorous or ironic effects (Attardo, 1994; Yus, 2016).

6.2.2 Analysis of Inter-Modal Incongruity

To address the IMI problem, there are typically two approaches: **1)** Always selecting one modality as the primary input. For instance, sentiment analysis, as well as humor and sarcasm detection have demonstrated the best performance when language is chosen as the primary input (M. K. Hasan et al., 2021; F. Ma, Zhang, & Sun, 2023; Rahman et al., 2020). However, *it cannot be guaranteed that one modality consistently and predominantly contributes to the final task*. Even though language is of utmost importance, it does not imply that every sample is dominated by language. **2)** Employing a weighted sum to adjust the input from each modality, thereby avoiding their equal contributions (Rahman et al., 2020; K. Yang, Xu, & Gao, 2020). Yet, *this approach overlooks the fact that some modalities can be more closely related or synchronized*. Fusing all modalities simultaneously without establishing a hierarchy may diminish the efficacy of fusion. Additionally, the practice of encoding the same input at both the word- and utterance-level for fusion has also been adopted to address incongruity (Y. Wu et al., 2021). However, *such an approach may introduce redundancy due to the repeated encoding of the same modality input*.

Furthermore, most research typically provides high-level examples to demonstrate how modalities may be misaligned. For instance, in a video clip from MUSTARD for sarcasm detection, a positive sentence—“You’re right, the party is fantastic,” is presented with a facial expression of “eye-rolling up” and drawn-out syllables on the word “fantastic” (Y. Wu et al., 2021). However, no evidence has been presented to illustrate how different modalities may mismatch at the latent level. Therefore, we perform an analysis using heatmaps to reveal how cross-attention highlights mismatched latent information. We conduct the following experiments on CMU-MOSEI. Note that approximately twice the number of examples presented were examined. We select the most representative cases for illustration, following previous work Rahman et al. (2020); Tsai, Bai, et al. (2019).

Exp 1. Investigating how source modality enhances target modality via cross-attention. We use the example of $V \rightarrow T$ (text attended by vision). Next, we aim to see how the combination of two modalities collectively affects the third:

Exp 2. Investigating how the salient parts of target modality are represented by self-attention with and without the combination of two source modalities. We use the example of $(A + V) \rightarrow T$ (text attended by cross-attention-fused audio-vision). Further, we hope to know how different source modalities individually affect the target:

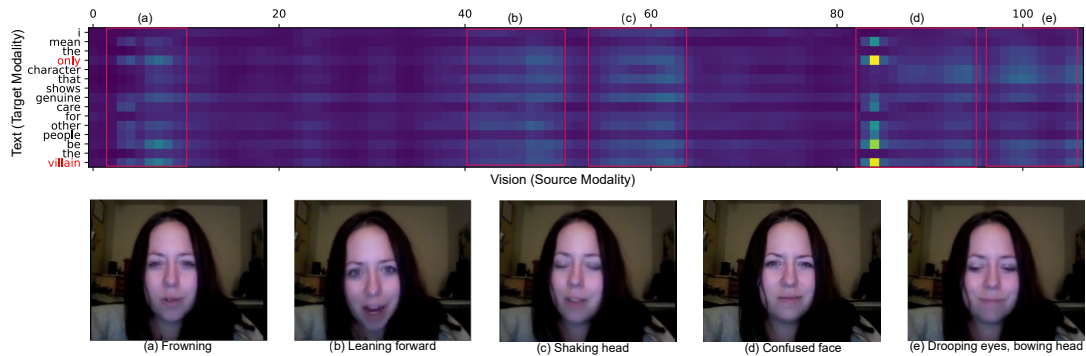
Exp 3. Investigating how the salient parts of target modality are represented by cross-attention when using different source modalities. We use the examples of $V \rightarrow T$ (text attended by vision) and $A \rightarrow T$ (text attended by audio).

The experimental setting is shown in Table 6.1, and the visualization is shown in Figure 6.1, 6.2, and 6.3.

Table 6.1: Experimental setting for multimodal incongruity analysis.

Exp.	Target modal	Source modal	Cross-attention	Self-attention
1	T	V	$V \rightarrow T$	/
2	T	$A + V$	$(A + V) \rightarrow T$	T
3	T	A or V	$A \rightarrow T, V \rightarrow T$	/

6.2.3 Results and Discussion

**Figure 6.1:** Exp 1 – Heatmap of highlighted hidden states using cross-attention on V and T .

Crossmodal Attention: but this movie kind of came off like somebody really **sad** travel video, i **don't** know

Text Self-attention: **but this movie** kind of came off like somebody really sad travel video, **i don't** know

Crossmodal Attention: and umm first of all, i, i actually was **never** a huge **fan** of the xx xxx tv series

Text Self-attention: **and umm first of all, i, i** actually was never a huge fan of the xx xxx tv series

Crossmodal Attention: i really **love** the **melodramatic** over the top crazy action sort of thing going on

Text Self-attention: **i really love the melodramatic over** the top crazy action sort of thing going on

Figure 6.2: Exp 2 – Heatmap of highlighted words using self-attention w/ and w/o cross-attention.

Crossmodal Attention ($V \rightarrow T$): i **love** him as an **actor** based on his looks, as an actual **talented** actor, **no**

Crossmodal Attention ($A \rightarrow T$): i **love** him as an actor based on his looks, as an actual **talented** actor, **no**

Crossmodal Attention ($V \rightarrow T$): like i said, **maybe my expectations were lowered because** of all of those reviews, **but** i really **didn't hate it**

Crossmodal Attention ($A \rightarrow T$): like **i** said, **maybe my expectations were lowered because** of all of those reviews, **but** i really **didn't hate it**

Figure 6.3: Exp 3 – Heatmap of highlighted words using different modalities in cross-attention.

Figure 6.1 shows the video frame (x-axis) and text words (y-axis). The salient affective information captured by cross-attention is highlighted in red boxes. It can be observed that the highlighted parts are due to obvious facial or behavior changes of the character in the video, such as frowning or shaking head. The cross-attention successfully highlights the meaningful words associated with a facial expression (e.g., “only”, “villain”).

In Figure 6.2, it can be noted that when fused with the combination of source modalities, T focuses more on the words related to emotion with less noise from other words. For example, when with cross-attention, the word “sad” is the most salient in the first sentence, yet much less important with self-attention. The same is true for the word “never” in the second sentence and the words “love” and “melodramatic” in the third sentence.

In Figure 6.3, we observe that when fused with different individual source modalities, the target modality T can be enhanced with disparate affective tendencies. When using V as the source modality, the words “love” and “talented” are the most highlighted in the first sentence, representing a positive meaning. When using A , however, “no” is the most focused word, showing negation is important. Similarly, “but” showing the turnaround in the second sentence is captured by V yet ignored by A , and the two draw attention to different parts. These phenomena demonstrate that different modalities may contain mismatched affective tendencies. The existence of IMI has been found by high-level inter-modal comparison (Desai, Chakraborty, & Akhtar, 2022), especially in sentiment analysis (Y. Li, Ishi, et al., 2019) and sarcasm detection (Y. Wu et al., 2021). This analysis indicates that incongruity also exists latently, resulting in hidden states in one modality being affected by the other.

Based on the above findings, we find that cross-attention does help multimodal fusion by aligning two modalities to highlight the salient affective information in the target modality with complementary information from the source. According to the attention mechanism (Vaswani et al., 2017), this process can be described as mapping the Query (from the target) to the Key (from the source) and obtain scores for the Value (from the source). However, such a process could malfunction if the modalities have mismatched affective tendencies, which leaves the IMI difficult to resolve at the latent level.

6.3 Incongruity-Aware Modality-Gated Fusion

To exploit the advantages of cross-attention while solving the above problems, we propose a new multimodal fusion approach: Hierarchical Crossmodal Transformer with Dynamic Modality Gating (HCT-DMG), which improves on existing methods in two aspects: **1)** Several previous approaches treated all modalities equally and fused them at every step, leaving incongruity in the fusion (Sahay, Okur, Kumar, & Nachman, 2020; Tsai, Bai, et al., 2019), while our HCT-DMG approach initially fuses the auxiliary modalities before integrating the primary modality in the final step. This strategy avoids excessive influence on the effect of the primary modality. **2)** Some prior work determined a primary modality based on the hierarchy of modalities used before modeling (Hazarika et al., 2020; Rahman et al., 2020). Such a practice is empirical and results in a fixed hierarchy. Thereby, the weighting pattern (e.g., $T \oplus W_1A \oplus W_2V$) cannot be changed during model training even though other hierarchies may

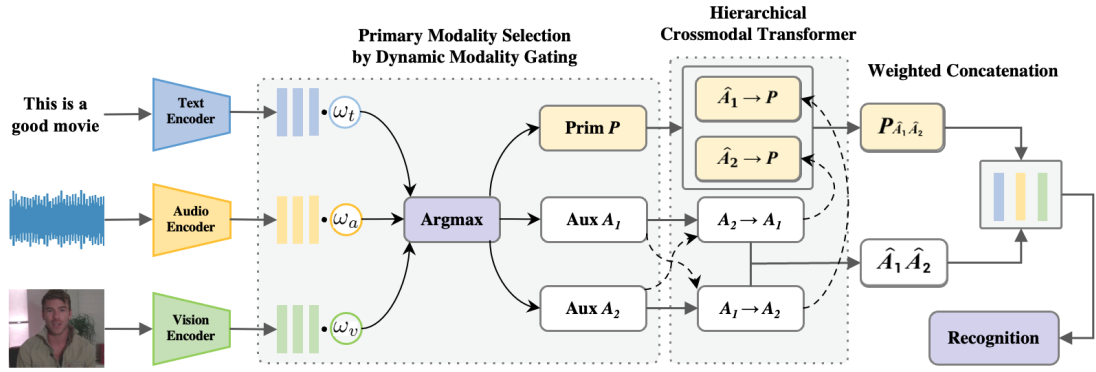


Figure 6.4: Architecture of HCT-DMG. Prim and Aux are short for primary and auxiliary. Dot lines denote crossmodal attention.

become better suited to the task. In contrast, HCT-DMG automatically selects and dynamically changes the primary modality in each training batch and constructs the hierarchy accordingly. Therefore, our proposed approach can *reduce redundancy*, *improve model reliability in incongruent scenarios* and enable the model to be *modality-agnostic*.

The architecture is shown as Figure 6.4. HCT-DMG is constructed based on three modalities: Text (T), Audio (A), and Vision (V), and consists of four components: feature encoder, dynamic modality gating, hierarchical crossmodal Transformer, and weighted concatenation. Note that HCT-DMG supports modalities not limited to T , A , and V , as the dynamic modality gating enables to construct the best hierarchy for any three inputs.

Feature Encoder. The input features are first fed into 1D Convolutional (Conv1D) networks to integrate local contexts and project the features into the same hidden dimension. Then the features are passed to the Gated Recurrent Unit (GRU) networks, which encode global contexts by updating their hidden states recurrently and model the sequential structure. We use two sets of input features: one uses the same conventional feature extractors as Tsai, Bai, et al. (2019); Tsai, Liang, Zadeh, Morency, and Salakhutdinov (2019) and Sahay et al. (2020) for comparison, while the other uses Large Pre-trained Models (LPM) for performance improvement, which will be described in Section 6.3.1.

Dynamic Modality Gating. DMG determines which modality should be the primary one by the trainable weight for each modality during training, rather than by manual selection. Specifically, each modality is assigned a trainable weight whose value is based on its contribution to the final task. The larger the contribution of a modality, the larger its weight value. The sum of all trainable weights equals to 1, and we allow the weights to be updated in every training batch to ensure that DMG can be well adapted to any input size by working dynamically. We will present the process of how to select the primary modality with DMG in Section 6.3.3.

Algorithm 1: Principle of HCT-DMG.

Input: Primary modality P ; Auxiliary modalities A_1 and A_2 ; Fixed modality index $Idx_P, Idx_{A_1}, Idx_{A_2} = 0, 1, 2$; Dynamic modality gating weight W initialized with $\text{nn.Param}(\text{softmax}([1, 1, 1]))$

```

1 repeat
2 for model input do
3   Sample Label, Text, Audio, Vision;
4    $W' \leftarrow \text{softmax}(W)$ ;
5    $P \leftarrow P * W'[Idx_P]$ ;
6    $A_1, A_2 \leftarrow A_1 * W'[Idx_{A_1}], A_2 * W'[Idx_{A_2}]$ ;
7    $Idx_{P'} \leftarrow \text{argmax}(W')$ ;
8   if  $Idx_{P'} == Idx_{A_1}$  then
9     |  $P, A_1 \leftarrow A_1, P$ 
10  end if
11  if  $Idx_{P'} == Idx_{A_2}$  then
12    |  $P, A_2 \leftarrow A_2, P$ 
13  end if
14   $Pred \leftarrow \text{HCT}(P, A_1, A_2)$ ;  $\text{update}(Pred, Label)$ ;
15 end for
16 until no model input

```

Hierarchical Crossmodal Transformer. As a variant of self-attention, cross-attention (J. Lu et al., 2019) transforms the signals from the source modality into a different set of Key-Value pairs to interact with the target modality (in multimodal fusion, cross-attention is usually referred to as crossmodal attention), which has proven useful in various domains (Rashed, Elsayed, & Schmidt-Thieme, 2022; S. Zhang, Chadwick, Ramos, & Bhattacharya, 2022). The crossmodal Transformer used here is the same as that of MulT (Tsai, Bai, et al., 2019), which is a deep stacking of several cross-attention blocks with layer normalization and positional embeddings. Unlike MulT, which has six crossmodal Transformers in the same step, we use two in the first step A_1 and A_2 to obtain enhanced auxiliary modalities \hat{A}_1 and \hat{A}_2 :

$$\hat{A}_1 = \text{CMT}(A_2 \rightarrow A_1) \quad (6.1)$$

$$\hat{A}_2 = \text{CMT}(A_1 \rightarrow A_2) \quad (6.2)$$

Then in the second step, another two crossmodal Transformers are used to yield the enhanced primary modality $\hat{P}_{\hat{A}_1}$ and $\hat{P}_{\hat{A}_2}$:

$$\hat{P}_{\hat{A}_1} = \text{CMT}(\hat{A}_1 \rightarrow P) \quad (6.3)$$

$$\hat{P}_{\hat{A}_2} = \text{CMT}(\hat{A}_2 \rightarrow P) \quad (6.4)$$

where P denotes the primary modality. The pseudo code of the working principle of HCT-DMG is shown in Algorithm 1.

Weighted Concatenation. After obtaining the enhanced $\hat{P}_{\hat{A}_1}$ and $\hat{P}_{\hat{A}_2}$, we concatenate them and use the self-attention to find its salient parts as the final primary representation $\hat{P}_{\hat{A}_1\hat{A}_2}$:

$$\hat{P}_{\hat{A}_1\hat{A}_2} = SA(\text{Concat} [\hat{P}_{\hat{A}_1}; \hat{P}_{\hat{A}_2}]) \quad (6.5)$$

At this point, crossmodal representations for every modality have been generated: \hat{A}_1 , \hat{A}_2 , and $\hat{P}_{\hat{A}_1\hat{A}_2}$. We concatenate them as the final multimodal representation Z :

$$Z = \text{Concat} [W_1\hat{A}_1; W_2\hat{A}_2; \hat{P}_{\hat{A}_1\hat{A}_2}] \quad (6.6)$$

where W_1 and W_2 are the weight matrices, which are learned by the model itself to control how much auxiliary information to extract.

6.3.1 Experimental Setup

Same as the analysis experiments, we use **CMU-MOSI**, **CMU-MOSEI**, **IEMOCAP**, **UR-FUNNY** and **MUSARD**. On MOSI and MOSEI, we evaluate the performances using the following metrics: 7-class accuracy (Acc7: sentiment score in the same scale as the labeled scores); binary accuracy (Acc2: positive/negative sentiment polarity); F1 score; Mean Absolute Error (MAE); and the correlation of the recognition results with ground truth (Corr). On IEMOCAP, we report the binary classification accuracy (one versus the others) and F1 score. On UR-FUNNY and MUSARD, we report the performance on binary classification accuracy (Acc2: humor/non-humor or sarcasm/non-sarcasm).

For MOSI, MOSEI, and IEMOCAP, we use the CMU-SDK (Zadeh, Liang, Poria, Vij, et al., 2018), which splits the datasets into training/validation/testing folds. For UR-FUNNY and MUSARD, the training, validation, and testing sets have already been split and provided. As described in Section 6.3, we use both conventional features and LPM features. The conventional features are obtained by using GloVe (Pennington et al., 2014), FACET (Littlewort et al., 2011), and COVAREP (Degottex, Kane, Drugman, Raitio, & Scherer, 2014) for T , V , and A , respectively. For the LPM, we use BERT (Devlin, Chang, Lee, & Toutanova, 2019) and WavLM (S. Chen et al., 2022) for T and A , respectively. We do not use LPM for V as they (e.g., CLIP (Radford et al., 2021)) did not show steady improvement as that for A and T .

Baselines

We perform a comparative study against our approach, considering four aspects: **1)** models using conventional features; **2)** models using LPM features; **3)** models using the same crossmodal Transformer as ours; **4)** models with similar sizes to ours. The baselines are as below:

Early Fusion LSTM (**EF-LSTM**) and Late Fusion LSTM (**LF-LSTM**) (Tsai, Liang, et al., 2019). Attention or Transformer-based fusion: **RAVEN** (Y. Wang et al., 2019), **MuIT** (Tsai, Bai, et al., 2019). Graph-based fusion: **Graph-MFN** (Zadeh, Liang, Poria, Cambria, & Morency, 2018). Low-rank-based fusion: **LMF** (Z. Liu et al., 2018). Cyclic translations-based fusion: **MCTN** (Pham, Liang, Manzini, Morency, & Póczos, 2019). Context-aware attention-based fusion: **CIA** (Chauhan, Akhtar, Ekbal, & Bhattacharyya, 2019). Multi-attention Recurrent-based fusion: **MARN** (Zadeh, Liang, Poria, Vij, et al., 2018). Temporal memory-based fusion: **MFN** (Zadeh, Liang, Mazumder, et al., 2018). Memory fusion by incorporating the information from the preceding context: **C-MFN**. Recurrent multiple stages-based fusion: **RMFN** (Liang et al., 2018). Low-rank Transformer-based fusion: **LMF-MuIT** (Sahay et al., 2020). Modality-invariant and -specific fusion using LPM: **MISA** (Hazarika et al., 2020). Integration of the preceding context and external knowledge: **HKT** (M. K. Hasan et al., 2021). Cascade residual autoencoder-based: **MMIN** (Zhao, Li, & Jin, 2021). Masked multihead attention-based: **TATE** (Zeng, Liu, & Zhou, 2022). Multimodal reconstruction-based: **MRAN** (W. Luo, Xu, & Lai, 2023). Modality translation-based: **TransM** (Z. Wang, Wan, & Wan, 2020) and **MTMSA** (Z. Liu, Zhou, Chu, Sun, & Meng, 2024). Graph network-based: **GCNet** (Lian, Chen, Sun, Liu, & Tao, 2023). Knowledge distillation-based: **CorrKD** (M. Li, Yang, Zhao, et al., 2024). Adversarial learning-based: **HRLF** (M. Li, Yang, Liu, et al., 2024). We also include several of the above-mentioned models enhanced by Connectionist Temporal Classification (CTC) (Tsai, Bai, et al., 2019).

Note that not all of the above baselines are employed for comparison across all five datasets for the following reasons:

1. As the recognition of fine-grained emotions can be largely improved by word-aligned data, we do not use the baselines with such alignment for a fair comparison on IEMO-CAP.
2. On MOSI, MOSEI, and IEMOCAP, the current State-Of-The-Art (SOTA) results are achieved by sophisticated approaches such as Self-MM (Yu, Xu, Yuan, & Wu, 2021), DMD (Y. Li, Wang, & Cui, 2023), CorrKD (M. Li, Yang, Zhao, et al., 2024), and MAG-XLNet (Rahman et al., 2020). However, their approaches fundamentally differ from the aforementioned baselines as they either generated additional unimodal labels using multimodal information and labels (i.e., ground-truth) or integrated multimodal inform-

ation into the Transformer by modifying its structure. Therefore, we exclude them from the comparison since all the models listed in the table concentrate exclusively on fusion methods without introducing extra training tasks or making modifications to the Transformers or pre-trained encoders.

3. Since humor and sarcasm detection are relatively new tasks in affective computing compared to sentiment analysis and emotion recognition, there are fewer baselines available. Therefore, we also include their variants that used different language models for comparison. We omit the size comparison as there are no comparable small models in the literature. For the HKT baseline, we use the performance without its proposed humor centric features because these features represent an additional modality, resulting in an unfair comparison with the others based on T , V , A . For the same reason in 2), we also exclude MAG-XLNet for comparison on UR-FUNNY and MUSTARD.

6.3.2 Results and Discussion

The comparison results are shown in Tables 6.2, 6.3, 6.4, and 6.5. Note that MulT and MISA are our reproduced results using their official code. † denotes models using LPM features. ◇ denotes models using the same crossmodal Transformer as ours. On all five datasets, it can be seen that HCT-DMG achieves better results on almost every metric than the baselines when using LPM features, showing the effectiveness of our proposed approach.

Moreover, when using conventional features, HCT-DMG improves every metric on MOSI and MOSEI, and almost every metric on IEMOCAP compared to the models of similar size. It also achieves highly competitive performance even compared to some LPM baselines on MOSI, UR-FUNNY, and MUSTARD. In addition, compared to LMF-MulT and MulT, which use the same crossmodal Transformer as ours, HCT-DMG still outperforms them, especially on MOSI and MOSEI by a large margin. These demonstrate the superiority of our proposed approach.

To the best of our knowledge, HKT is the SOTA model on humor and sarcasm detection in the literature thus far. Our result on UR-FUNNY does not surpass theirs, but we achieve higher result on MUSTARD, demonstrating the competitiveness of our model compared to the SOTA. Note that HKT was developed specifically for experimentation on these two datasets with a punchline detection module, its effectiveness is unlikely to generalize to MOSI, MOSEI, and IEMOCAP for sentiment analysis and emotion recognition. On the contrary, HCT-DMG performs well on all five datasets, demonstrating that our proposed approach can address the IMI problem not only in general incongruous scenarios (humor and sarcasm) but also in complex ones (sentiment and emotion). Therefore, the generalizability of HCT-DMG is well-established, allowing for its extension to various MIP tasks.

Furthermore, as the cross-attention results in Exp. 2 in Section 6.2 are produced by using HCT-DMG, Figure 6.2 clearly shows that our approach strengthened the salient affective parts.

Table 6.2: Comparison results of sentiment analysis on MOSI.

Model	CMU-MOSI				
	Acc7 \uparrow	Acc2 \uparrow	F1 \uparrow	Corr \uparrow	MAE \downarrow
EF-LSTM	33.7	75.3	75.2	0.608	1.023
RAVEN	33.2	78.0	76.6	0.691	0.915
MCTN	35.6	79.3	79.1	0.676	0.909
CTC+EF-LSTM	31.0	73.6	74.5	0.542	1.078
CTC+RAVEN	31.7	72.7	73.1	0.544	1.076
CTC+MCTN	32.7	75.9	76.4	0.613	0.991
MARN	34.7	77.1	77.0	0.625	0.968
MFN	34.1	77.4	77.3	0.632	0.965
RMFN	38.3	78.4	78.0	0.681	0.922
LMF	32.8	76.4	75.7	0.668	0.912
CIA	38.9	79.8	79.5	0.689	0.914
MTMSA	-	84.9	-	-	-
MMIN	-	82.3	-	-	-
TATE	-	84.9	-	-	-
MRAN	-	83.9	-	-	-
TransM	-	-	82.6	-	-
GCNet	-	-	83.2	-	-
CorrKD	-	-	83.9	-	-
HRLF	-	-	84.2	-	-
MISA \dagger	41.4	81.9	81.8	0.762	0.810
MuIT \diamond (1.07M)	34.3	80.3	80.4	0.645	1.008
LMF-MuIT \diamond (0.86M)	34.0	78.5	78.5	0.681	0.957
LF-LSTM (1.24M)	33.7	77.6	77.8	0.624	0.988
HCT-DMG					
<i>Conven.</i> (0.78M)	39.4	82.5	82.5	0.710	0.881
<i>LPM</i> \dagger (0.83M)	41.8	85.1	84.8	0.732	0.855

Finally, we present some hard samples in Table 6.6, where incongruity exists. It can be observed that MuIT (using the same crossmodal Transformer as ours yet repeatedly fusing every modalities) fails to handle these hard samples (videos available²) producing results that contradict ground truths. In contrast, our approach can recognize true sentiments with very close scores.

In summary, the effectiveness of our proposed HCT-DMG has been demonstrated from different angles by the experimental results, the heatmap visualization, and the recognition scores of several hard samples.

2. <https://sites.google.com/view/tacls submission>

Table 6.3: Comparison results of emotion recognition on MOSEI.

Model	CMU-MOSEI				
	Acc7 \uparrow	Acc2 \uparrow	F1 \uparrow	Corr \uparrow	MAE \downarrow
EF-LSTM	47.4	78.2	77.9	0.642	0.616
RAVEN	50.0	79.1	79.5	0.662	0.614
MCTN	49.6	79.8	80.6	0.670	0.609
CTC+EF-LSTM	46.3	76.1	75.9	0.585	0.680
CTC+RAVEN	45.5	75.4	75.7	0.599	0.664
CTC+MCTN	48.2	79.3	79.7	0.645	0.631
LMF	48.0	82.0	82.1	0.677	0.623
Graph-MFN	45.0	76.9	77.0	0.540	0.710
CIA	50.1	80.4	78.2	0.590	0.680
TransM	-	-	81.5	-	-
GCNet	-	-	82.4	-	-
CorrKD	-	-	82.2	-	-
HRLF	-	-	82.9	-	-
MISA \dagger	51.8	84.2	84.0	0.724	0.568
MuIT \diamond (1.07M)	50.4	80.7	80.6	0.677	0.617
LMF-MuIT \diamond (0.86M)	49.3	80.8	81.3	0.668	0.620
LF-LSTM (1.24M)	48.8	77.5	78.2	0.656	0.624
HCT-DMG					
<i>Conven.</i> (0.78M)	50.6	81.6	81.9	0.691	0.593
<i>LPM</i> \dagger (0.83M)	53.2	84.2	84.0	0.752	0.535

Table 6.4: Comparison results of emotion recognition on IEMOCAP.

Model	IEMOCAP							
	Happy		Sad		Angry		Neutral	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1
CTC+EF-LSTM	76.2	75.7	70.2	70.5	72.7	67.1	58.1	57.4
CTC+RAVEN	77.0	76.8	67.6	65.6	65.0	64.1	62.0	59.5
CTC+MCTN	80.5	77.5	72.0	71.7	64.9	65.6	49.4	49.3
MuIT \diamond (1.07M)	85.6	79.0	79.4	70.3	75.8	65.4	59.5	44.7
LMF-MuIT \diamond (0.86M)	85.6	79.0	79.4	70.3	75.8	65.4	59.2	44.0
LF-LSTM (1.24M)	72.5	71.8	72.9	70.4	68.6	67.9	59.6	56.2
HCT-DMG								
<i>Conven.</i> (0.78M)	85.6	79.0	79.4	70.3	75.8	65.4	61.0	50.5
<i>LPM</i> \dagger (0.83M)	87.1	81.6	82.4	73.2	79.0	68.8	63.2	60.3

6.3.3 Analysis and Discussion

To verify that our approach alleviates the IMI issue as well as to demonstrate how DMG dynamically changes the primary modality, we conduct some further studies. To showcase the advantages of our approach and for brevity, we present the analysis results using only conventional features.

Table 6.5: Comparison results of humor and sarcasm detection on UR-FUNNY and MUSTARD.

Model	UR-FUNNY	MUSTARD
	Acc2↑	Acc2↑
C-MFN (GloVe)	65.23	-
C-MFN (ALBERT) [†]	61.72	-
MISA (GloVe)	68.60	-
MISA (BERT) [†]	69.62	66.18
MISA (BERT) [†] [punchline only]	70.61	-
MISA (ALBERT) [†]	69.82	66.18
HKT [†]	76.36	75.00

HCT-DMG		
<i>Conven.</i>	68.54	73.29
<i>LPM</i> [†]	75.09	76.62

Table 6.6: Examples containing incongruity from CMU-MOSI.

#	Spoken words + acoustic and visual behaviors	Ground truth	MuT	Ours
1	“And that’s why I was not excited about the fourth one.” + Uninterested tone and facial expression	-1.4	1.185	-1.416
2	“I give Shrek Forever After directed by Mike Mitchell a grade of B minus.” + Smile face	1.0	-0.576	0.959
3	“Um in general um, the little kids seemed to like it that were in there.” + Skeptical tone and facial expression	0.8	-1.151	0.700
4	“I honestly want the aliens to win.” + Negative tone (somewhat disdainful)	-1.6	0.995	-1.906

Alleviation of the IMI Problem

MuT: And that’s why I was not **excited** about the fourth one.

Ours: And that’s **why** I was **not** excited about the **fourth** one.

MuT: i **love** him as an actor based on his looks, as an actual **talented** actor, **no**

Ours: **i love him as an actor** based on his looks, as an actual **talented** actor, no

Figure 6.5: Comparison examples of IMI.

First, as shown in Figure 6.3 in Section 6.2, the affective words are enhanced by the auxiliary modalities. However, *V* focuses more on positive words while *A* highlights negation the most, which likely changes affective tendency. We re-implemented MuT and extracted the attention of its enhanced *T* in the final stage to compare with ours. In the first example of Figure 6.5, it can be seen that MuT is confused by the IMI, resulting in the negative word being barely focused on, which yields a completely opposite affective tendency to the ground truth. In the second example, the positive and negative words are treated equally by MuT, leaving the IMI unsolved. These occur because MuT fuses *V* and *A* with *T* at the same level, and simply

concatenates two enhanced T . In contrast, our approach gives little attention to the word “no”, showing that the IMI is resolved at the latent level as the hidden states of “no” are barely encoded by the cross-attention. Together with the examples in Table 6.6, it is observed that our approach can greatly alleviate the IMI problem for these examples.

Automatic Modality Selection by DMG

As the DMG automatically selects the primary modality by adjusting the weight for each modality, we show how the weights vary during training using CMU-MOSI. The weight of a modality denotes the confidence that this modality is selected as the primary one. Figure 6.6 shows how weights of the modalities vary in the first epoch. It can be noted that T is not the primary one at the beginning but gradually dominates after batch 60. Figure 6.7 shows the variation in the average weight of each modality with epoch. It can be seen that T indeed dominates, and the weight distribution starts to converge at around epoch 40. Meanwhile, V gradually surpasses A with epoch although the opposite is true in the first epoch in Figure 6.6. It can be inferred that DMG can dynamically change the weights as the training goes on, demonstrating the efficacy of our approach.



Figure 6.6: Weight variation in the first epoch on MOSI dataset.

Furthermore, since HCT-DMG works without the need to understand which modalities are used, its utility can be expected to expand beyond more general scenarios with different inputs.

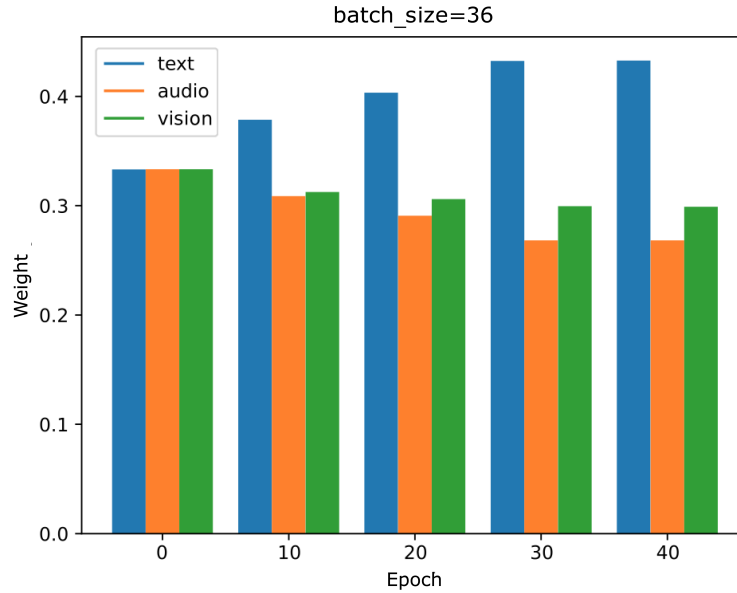


Figure 6.7: Variation of average weight with epoch on MOSI dataset.

Ablation Study on Removal of DMG

Although DMG automatically selects T as the primary modality, we hope to see if this phenomenon really brings the best performance compared to prior work whose hierarchies are fixed. Thus, we remove DMG and construct an HCT model with three hierarchies, in each of which either T , A , or V was manually selected as the primary modality respectively for performance comparison. As shown in Table 6.7, selecting T as the primary modality achieves the best performance on every metric on MOSI and MOSEI and on most metrics on IEMOCAP, verifying the rationality and efficacy of the automatic modality selection by DMG.

Moreover, although the results of selecting T as the primary modality are the best among the three, overall scores on all three datasets are slightly lower than HCT-DMG’s in Tables 6.2, 6.3, 6.4, and 6.5. This phenomenon is reasonable because the primary modality keeps changing during training, even though T is primary overall. Such a dynamic property cannot be encoded by the manual selection of the primary modality whereas it can be encoded by HCT-DMG and thus yields better results. This is consistent with the weight variation observed in Figure 6.6 and 6.7, which reaffirms the validity of our approach. In addition, the results on IEMOCAP, where no modality consistently dominates, align with their close weights (all slightly above 0.33 by the experiment in Section 6.3.3 using IEMOCAP).

We omit UR-FUNNY and MUSTARD for brevity as the prior work by HKT has explored the contribution, where each modality was individually used, and the other two were removed (M. K. Hasan et al., 2021). They observed that in humor detection, the contribution of each modality followed the order of $T > A > V$, while in sarcasm detection, the order is $T = A > V$.

Table 6.7: Performance comparison by manually selecting different primary modalities with convention features.

Primary modality	CMU-MOSI				
	Acc-7↑	Acc-2↑	F1↑	Corr↑	MAE↓
<i>T</i>	38.9	82.5	82.6	0.717	0.859
<i>A</i>	37.5	81.3	81.3	0.705	0.883
<i>V</i>	38.3	80.9	81.0	0.679	0.909

Primary modality	CMU-MOSEI				
	Acc-7↑	Acc-2↑	F1↑	Corr↑	MAE↓
<i>T</i>	50.1	81.8	81.9	0.685	0.601
<i>A</i>	47.5	79.6	80.3	0.650	0.644
<i>V</i>	48.7	80.8	81.0	0.659	0.633

Primary modality	IEMOCAP							
	Happy		Sad		Angry		Neutral	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1
<i>T</i>	85.6	79.4	79.5	70.6	75.8	65.4	59.6	55.6
<i>A</i>	85.4	79.4	78.8	70.4	75.4	65.5	59.3	52.4
<i>V</i>	85.6	79.2	79.4	70.3	75.9	65.6	60.4	53.3

6.3.4 Integration with LLMs

To further validate our approach, we explore the performance using LLMs. Given the necessity of a sufficient amount of data to fine-tune an LLM into a multimodal LLM, we combine MOSI, MOSEI, and IEMOCAP. To align these datasets, we unify their emotion labels into three categories: *positive*, *negative*, and *neutral*. These datasets are trained together, with separate evaluations conducted for each.

We use *LlaMA-3.1-8b* (Dubey et al., 2024) as the LLM and employ a single-layer linear projection to align the multimodal inputs with the text feature space, following the setup of MiniGPT-4 (D. Zhu, Chen, Shen, Li, & Elhoseiny, 2024). Furthermore, to perform end-to-end emotion-aware feature extraction for multimodal alignment, we use *emotion2vec plus large* as the speech encoder and HSemotion as the visual encoder. Emotion2vec (Z. Ma et al., 2024) is a self-supervised speech emotion representation model designed to extract universal emotional features from speech signals. It utilizes self-supervised online distillation, combining utterance-level and frame-level losses during pre-training on open-source unlabeled emotion data. HSemotion (Savchenko, 2022) is a high-speed and lightweight facial emotion recognition model. It employs multi-task learning to simultaneously predict basic emotions, valence-arousal values, and expression intensity from images or video frames. We employ Emotion2vec and HSemotion to extract emotion-aware speech and visual representations, aiming to enhance the emotional capabilities of the LLM. Due to resource constraints, the model is trained with a batch size of 2 for 5 epochs. The architecture is shown in Figure 6.8.

The task command, emotional sentence, and available emotion options are integrated into a single text input to prompt the LLM: “Please analyze the emotion of the given sentence and choose the emotion from {Positive, Neutral, Negative}”. Since LLMs are primarily trained on textual data, we designate text as the primary modality to leverage their strong reasoning abilities in this domain (in other words, assigning audio or visual input as the primary modality severely impairs their reasoning performance).

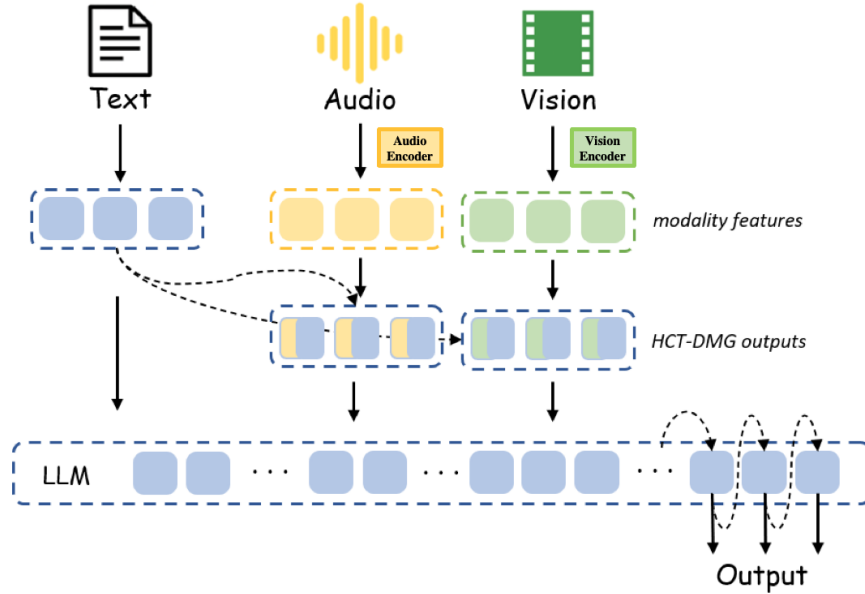


Figure 6.8: LLM integrated with HCT-DMG (LLM-HCT).

Considering that fine-tuning a pure text-based LLM model to a multimodal model requires a large amount of data, we concatenate three datasets for joint training, but evaluate them separately. Specifically, we standardize the labels of the MOSI, MOSEI, and IEMOCAP datasets into three categories: positive, neutral, and negative, as shown in Table 6.8. Audio and visual inputs are encoded into 1024 and 1280 dimensions, respectively, and then projected into 4096 dimensions to align with the text feature dimension of *LlaMA-3.1-8b*.

Table 6.8: The data distribution of the dataset used for training the LLM.

Sentiment	MOSI			MOSEI			IEMOCAP			Total
	Train	Valid	Test	Train	Valid	Test	Train	Valid	Test	
Positive	680	124	277	8,062	932	2,289	1,393	180	176	14,113
Neutral	54	13	30	3,544	433	1,026	1,331	161	161	6,753
Negative	552	92	379	4,572	507	1,353	3,194	400	400	11,449
Total	1,286	229	686	16,178	1,872	4,668	5,918	741	737	32,315

We evaluate our LLM-HCT in comparison with two baseline strategies:

- **LLM-Concat:** Features from the audio and visual modalities are concatenated to those from the text modality.

Table 6.9: Comparison results using LLMs.

Setting	Audio	Visual	MOSI		MOSEI		IEMOCAP	
			Acc	F1	Acc	F1	Acc	F1
LLM-Concat	✓	✗	78.51	78.51	72.17	72.17	76.64	76.64
LLM-Concat	✓	✓	77.97	77.97	72.66	72.66	78.56	78.56
LLM-Cross	✓	✗	78.26	78.26	71.94	71.94	76.60	76.60
LLM-Cross	✓	✓	77.43	77.43	72.06	72.06	77.91	77.91
LLM-HCT	✓	✗	80.17	80.17	72.49	72.49	77.42	77.42
LLM-HCT	✓	✓	79.70	79.70	73.05	73.05	78.93	78.93

- **LLM-MuIT**: Each pair among the three modalities undergoes cross-modal attention, and the resulting outputs are then concatenated together (same as the fusion strategy of MuIT (Tsai, Bai, et al., 2019)).
- **LLM-HCT**: The three modalities are combined using our proposed HCT-DMG (as in Figure 6.8).

As shown in Table 6.9, our proposed scheme still outperforms MuIT within the LLM framework and achieves the best performance among the three approaches. Interestingly, MuIT leads to a performance drop, even compared to simple concatenation, possibly due to the added complexity of aligning modalities through cross-modal attention. The repeated fusion of visual and audio information may hinder the LLM from effectively learning relevant information from the textual modality. In contrast, LLM-HCT designates text as the primary modality and hierarchically integrates information from the other two modalities into it, minimizing divergence from the original text feature space and enabling more effective adaptation of visual and audio features to the text modality. Finally, it is noticed that with the visual modality, the results are worse than those without visual modality on MOSI. We believe that it is likely due to two reasons: **1)** Audio and text typically align well, particularly in tasks involving emotion, sentiment, or intent. Since the LLM is pre-trained solely on text, aligning visual features with text is more challenging. **2)** Introducing a third modality increases model complexity (particularly because visual features often lack consistency with textual and audio features, as noted in Reason 1) and increases data requirements. Due to the limited size of MOSI, the model is susceptible to overfitting on the training data.

6.3.5 Summary

In this section, we analyzed cross-attention-based multimodal fusion and propose a Hierarchical Crossmodal Transformer with Dynamic Modality Gating for incongruity-aware multimodal affect recognition. **1)** We demonstrate the existence of inter-modal incongruity at the latent level due to cross-attention. Specifically, we show that cross-attention can help to capture affective information across modalities and enhance salient parts in the target modality, but it can also induce mismatched affective tendencies from different modalities. **2)** We proposed HCT-DMG, which dynamically selects the primary modality during training. This model requires fewer fusion operations and does not repeatedly fuse a single modality, reducing the parameters to approximately 0.8M while largely outperforming existing models of similar size. **3)** We further analyzed the mechanism and feasibility of automatic modality selection by DMG and show that the selection process supports the primacy of text in prior MIP studies, adding new insights to the literature. **4)** We tested the performance of HCT-DMG on five datasets for sentiment analysis, emotion recognition, humor and sarcasm detection. HCT-DMG achieves remarkable results, consistently outperforming almost all of the baseline models across all tasks. It highlights the versatility and effectiveness of our proposed approach in handling various multimodal affect recognition tasks. **5)** We expanded the experiments using LLMs, integrating multimodal features into *LlaMA-3.1-8b*. The results further demonstrates the generalizability and effectiveness of our proposed method.

6.4 ASR Error-Robust Modality-Gated Fusion

Since ASR errors bring noise to the text modality, it is expected that they could result in incongruity between the text and other modalities. In this section, we propose an ASR Error-Robust Modality-Gated Fusion that combines our ASR error correction techniques from Chapter 5 and the incongruity-aware modality-gated fusion from Section 6.3.

6.4.1 Experimental Setup

Settings

We experimented on both bimodal SER (text + audio) and trimodal (text + audio + vision) for a comprehensive SER evaluation. For the bimodal scenario, we use the same datasets and experimental settings as the benchmarking experiment in Section 4.3.3. For the trimodal scenario, we follow the same settings as in the previous section (Section 6.3).

Incongruity from ASR Errors

<i>Transcription with attention heatmap</i>	<i>Prediction</i>	<i>Prediction</i>	<i>Prediction</i>
	<i>text only</i>	<i>audio only</i>	<i>text + audio</i>
<i>Groundtruth:</i> No it's nothing like that. It's-.	Sad	Sad	Sad
<i>hubertlarge:</i> o knows what thin mit tanis	Neu	Sad	Ang
<i>w2v2100:</i> thou no something like that's	Sad	Sad	Neu
<i>whispermedium:</i> I'm doing something like that	Neu	Sad	Neu

Table 6.10: Incongruity brought by ASR errors. The groundtruth label is *Sad*.

First of all, we present a case study to show the incongruity brought by ASR errors. As shown in Table 6.10, the ASR error in the text modality causes incongruity with the audio modality. We compute the attention score and create the attention heatmap during the attention fusion of text and audio for SER. The attention heatmap successfully captures the salient word “nothing” in the ground truth transcription but fails to do so in the ASR transcriptions due to recognition errors. Even when either the text or audio modality alone correctly identifies the emotion from the ASR output, their fusion results in incorrect predictions. This case study underscores the misalignment between the text and audio modalities introduced by ASR errors, which can degrade performance in multimodal SER tasks.

6.4.2 Proposed Framework

To address both the WER and incongruity issues, we propose an ASR Error-Robust Modality-Gated Fusion framework involving ASR error correction and modality-gated fusion, as illustrated in Figure 6.9 (left side). This framework integrates insights from previous sections, specifically the LLM-based ASR error correction (Alpaca prompt module) and the S2S-based ASR error correction from Chapter 5, as well as the modality-gated fusion from Section 6.3.

ASR Error Correction

We combine the transcriptions from all used ASR models, forming N-best (N=11) hypotheses. Compared to using N-best hypotheses directly from a single ASR model, which offer limited diversity, such a combination offers a more diverse range of transcription results. Furthermore, this approach facilitates the comparison of the SER performance with each of the ten ASR models, as depicted in Table 4.8.

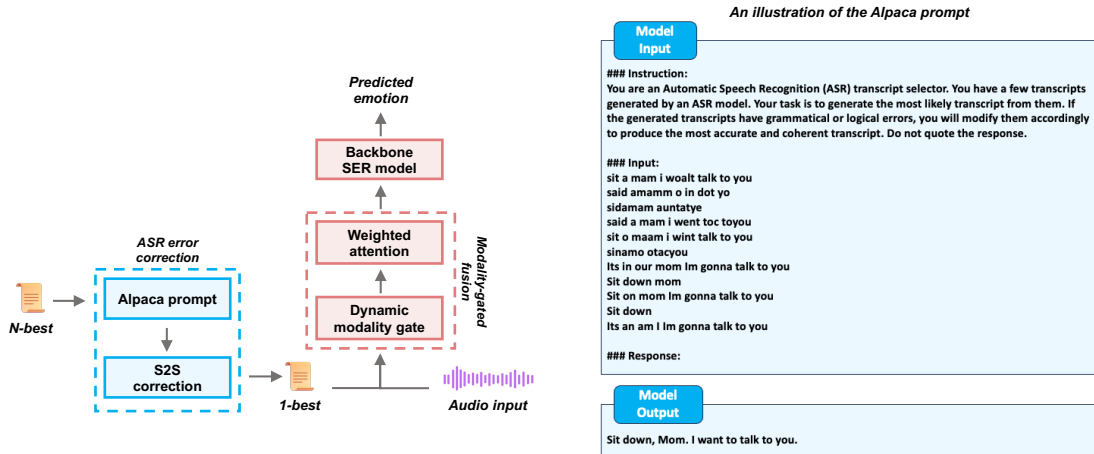


Figure 6.9: Proposed ASR error-robust modality-Gated fusion framework. **Blue:** ASR error correction; frozen. **Pink:** Modality-gated fusion; trainable.

Then, the ASR error correction process consists of a two-stage process. First, we utilize the Alpaca prompt (Taori et al., 2023), which guides an LLM to generate the most likely transcription (i.e., 1-best hypothesis) from the N-best hypotheses and correct any grammatical or logical errors. An illustration of the Alpaca prompt is depicted on the right of Figure 6.9. We employ InstructGPT, which has demonstrated effectiveness in LLM-based ASR error correction (Radhakrishnan et al., 2023).

In the second stage, we further refine the 1-best transcription by utilizing our proposed S2S ASR error correction model in Chapter 5, which was pre-trained on the *Whisper-tiny* output y' (i.e., 1-best hypothesis) of Common Voice to recover the corrected sequence y (i.e., gold annotation), with an optimizable parameter θ :

$$y = \underset{y}{\operatorname{argmax}} P(y|y', \theta) \quad (6.7)$$

Modality-Gated Fusion

Based on the quality-enhanced transcription, we adapt the trimodal incongruity-aware fusion method proposed in Section 6.3 to our bimodal inputs, considering potential distortion of original emotion in the spoken content due to ASR errors, which may introduce incongruence between emotions in audio and text.

Specifically, each modality is initially assigned an equal trainable weight. These weights dynamically adjust during training based on the respective contribution of each task, with larger contributions resulting in higher values. The sum of the weights always equals 1, and they are updated in every training batch to ensure dynamic adaptation to any input batch size. Subsequently, the weights are multiplied by corresponding modality features to perform weighted cross-attention fusion. The pseudo-code is shown in Algorithm 2, where H represents the

hidden representation as input to the backbone SER model. *CrossAttn* and *SelfAttn* denote multihead attention, the former with the first argument as query and the second as key and value, and the latter with the argument as query, key, and value. The head number is set to eight.

Algorithm 2: Modality-gated fusion.

Input: Audio modality A ; Text modality T ; Dynamic modality gating weight W^1 initialized with $\text{nn.Param(torch.tensor([1, 1]))}$; Feature concatenation weight W^2 initialized with $\text{nn.Param(torch.tensor([1, 1, 1]))}$

```

1 for model input do
2   Sample  $A, T$ ;
3    $W^1 \leftarrow \text{Softmax}(W^1)$ ;
4    $W^2 \leftarrow \text{Softmax}(W^2)$ ;
5    $W_A^1 \leftarrow W^1[0], W_T^1 \leftarrow W^1[1]$ ;
6    $W_A^2 \leftarrow W^2[0], W_T^2 \leftarrow W^2[1], W_{AT}^2 \leftarrow W^2[2]$ ,
7   if  $\text{argmax}(W^1) == 0$  then
8      $A' = W_A^1 * A$ ;
9      $T' = W_T^1 * \text{CrossAttn}(A, T)$ ;
10  else
11     $A' = W_A^1 * \text{CrossAttn}(T, A)$ ;
12     $T' = W_T^1 * T$ ;
13  end if
14   $H = \text{SelfAttn}(\text{Concat}[A', T'])$ ;
15   $H' = \text{Concat}[W_A^2 * A, W_T^2 * T, W_{AT}^2 * H]$ ;
16 end for

```

Finally, we apply the same backbone SER model, keeping all other settings unchanged (e.g., the text and audio encoders) as in Section 4.3.3. We believe that this combination of dynamic modality gate and weighted attention can effectively mitigate the negative impact of ASR errors by consistently focusing on the most salient parts in both inputs for SER.

6.4.3 Results

The results of the bimodal scenario are presented in Table 6.11. We include only the MAE on CMU-MOSI and the average CCC on MSP-Podcast for brevity. Our approach further reduces WER and improves SER performance, demonstrating its effectiveness. While it does not surpass the best transcription in terms of MAE on CMU-MOSI, the values are very close. Moreover, in line with the findings in J. He, Shi, Li, and Toda (2024), our results confirm that employing proper ASR error-robust approaches can exceed SER performance based on ground-truth text using ASR transcription. This phenomenon likely occurs because the LLM/S2S models use context to rewrite the transcript in a way that better aligns with the projected emotion, causing certain misrecognized words to be interpreted as contributing positively to the ground-truth emotions, rather than indicating that ASR errors are preferable in SER.

Table 6.11: Performance comparison. \uparrow : higher the better. \downarrow : lower the better.

	IEMOCAP		CMU-MOSI		MSP-Podcast	
	WER \downarrow	Acc4 \uparrow	WER \downarrow	MAE \downarrow	WER \downarrow	CCC \uparrow
best trans	12.31	74.66	18.69	0.8558	15.02	0.616
ground truth	0.00	74.32	0.00	0.8902	0.00	0.613
ours	10.12	76.66	17.00	0.8557	12.85	0.618

6.4.4 Summary

In this section, we proposed an ASR Error-Robust Modality-Gated Fusion that integrates ASR error correction and modality-gated fusion. This framework demonstrated superior performance compared to baseline models in our experimental results and holds potential for various tasks where ASR transcription serves as the text source.

6.5 Chapter Summary

Building on the foundations laid in earlier chapters, this chapter examines cross-modal incongruity in SER, arising from either modality misalignment or imperfections in ASR. Such incongruity may introduce conflicting emotional signals and undermine SER accuracy. To tackle this, we incorporate our proposed ASR error correction methods and present novel incongruity-aware fusion models to enhance the robustness of multimodal SER systems.

We begin by analyzing cross-modal incongruity in multimodal SER inputs, as detailed in Section 6.2. To mitigate this issue, Section 6.3 proposes a hierarchical attention-based fusion framework that employs an incongruity-aware modality-gated mechanism to reduce the effects of conflicting emotional cues. Additionally, in Section 6.4, we augment this framework by integrating ASR error correction into the modality-gated fusion process, resulting in an ASR error-robust fusion approach aimed at minimizing incongruity caused by ASR errors and improving overall SER outcomes.

Semi-Supervised Labeling for Effort-Saving SER

7.1 Introduction

After exploring ASR models, correcting ASR errors, and integrating ASR outputs, we aim to address another issue in SER: significant human efforts in emotion labeling, through the insights about ASR we have obtained so far. Furthermore, we also aim to evaluate our approach in another speech-based cognitive task: Alzheimer’s dementia detection, for generalizability.

Speech classification tasks, particularly those involving cognitive states, are valuable for a wide range of applications, including human-computer interaction (Deng et al., 2004), health monitoring (Fleury, Noury, Vacher, Glasson, & Seri, 2008), and clinical diagnosis (Moell, Aronsson, Östberg, & Beskow, 2025). However, human annotation for these tasks is often expensive, time-consuming, and requires extensive subjective assessment. This data scarcity issue hinders the progress of these classifications in real-world applications. To address this challenge, Semi-Supervised Learning (Semi-SL) offers a promising approach by leveraging both limited labeled data and a larger amount of unlabeled data for various classification tasks, such as hate speech detection (d’Sa, Illina, Fohr, Klakow, & Ruitter, 2020), emotion recognition and sound event detection (Z. Zhang & Schuller, 2012), as well as sleepiness detection and gender classification (Z. Zhang, Deng, & Schuller, 2013).

Generally, Semi-SL methods can be categorized into two main types: generating *reliable pseudo-labels* and building *reliable models* using limited data. For example, pseudo-labeling involves creating labels for unlabeled data based on predictions from an iteratively trained model (Z. Zhu & Sato, 2021). Consistency regularization ensures that the model produces consistent predictions for augmented versions of the same input (e.g., with added noise) (K. Lu, Foo, Teh, Tran, & Chandrasekhar, 2019). These two types are often complementary and are commonly used together in existing Semi-SL frameworks (T. Feng & Narayanan, 2022). Moreover, the advancements in self-supervised learning have further improved Semi-SL by utilizing large amounts of unlabeled data to pre-train Semi-SL models (Lai, Chuang, Lee, Li, & Glass, 2021; B. Zhang et al., 2022).

Among the literature, the most relevant studies to our work focus on pseudo-label generation. d'Sa et al. (2020) used label propagation, transducing labels from labeled data to unlabeled data with probabilistic transitions for hate speech classification. They found that semi-supervised learning based on label propagation improves hate speech classification in extremely low-resource scenarios, but the performance gains diminish as the amount of labeled data increases. Z. Zhang and Schuller (2012) incorporated unlabeled data with high confidence levels (i.e., classification probabilities generated by the trained classifier) into the training set and resampled the original labeled data for sound event classification. They also proposed dividing acoustic features into two views, selecting high-confidence instances in each view, and aggregating them with their predictions into the initial training sets per iteration (Z. Zhang et al., 2013). Z. Zhu and Sato (2021) applied noisy student training (Xie, Luong, Hovy, & Le, 2020) to emotion recognition, using a teacher model trained on labeled data to infer soft labels for unlabeled data. T. Feng and Narayanan (2022) proposed incorporating federated learning, utilizing both labeled and unlabeled data at each local client in a multi-view pseudo-labeling approach.

Despite these advances, selecting high-confidence pseudo-labeled data remains challenging. In this chapter:

- We propose a novel Semi-SL framework, integrating multi-view pseudo-labeling that leverages both acoustic (Chapter 3) and linguistic characteristics (Chapter 4) to select the most confident data for model training.
- We employ Fréchet Audio Distance (FAD) as a reference-free method to cluster unlabeled data based on acoustic similarity.
- We use task-specific prompts from Chapter 5 to predict labels from ASR transcripts, leveraging insights from acoustics, linguistics, and psychology.
- We examine multiple fusion methods, including the modality-gated fusion from Chapter 6 in the context of Semi-SL to build the bimodal classifier.

7.2 Proposed Model

As illustrated in Figure 7.1, the Semi-SL framework with our proposed multi-view pseudo-labeling consists of two paths: acoustic and linguistic. The acoustic path utilizes the similarity between labeled and unlabeled data based on diverse audio embeddings, while the linguistic path employs LLMs to predict class labels from ASR transcripts using task-specific knowledge. If the generated pseudo-labels from both paths align, we consider them as high-confidence data for training a bimodal classifier. Otherwise, the data are treated as low-confidence and will be further predicted using the trained bimodal classifier. The semi-supervised training of the bimodal classifier will iterate until a predefined criterion is met.

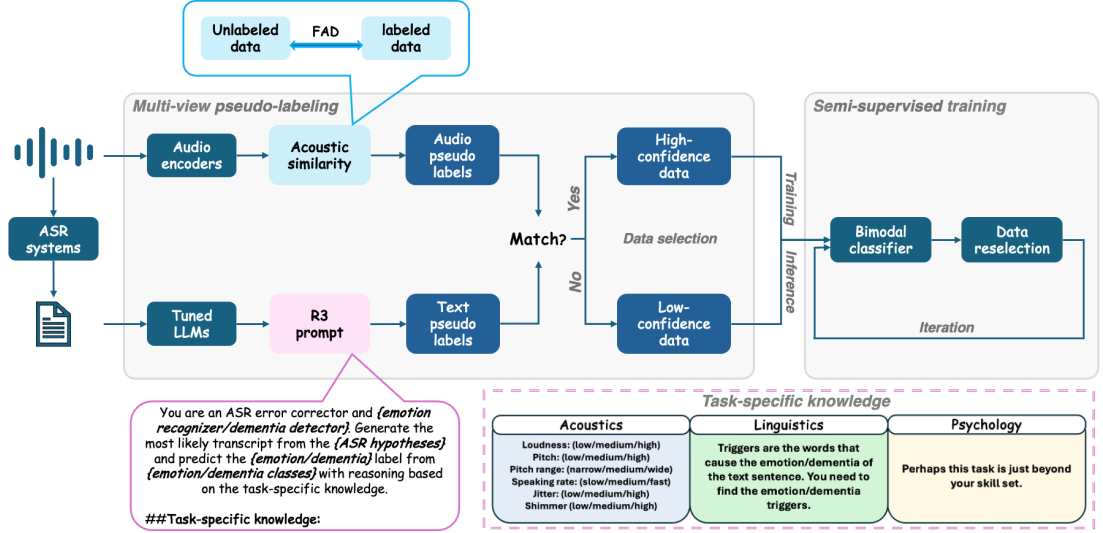


Figure 7.1: Framework of our proposed semi-supervised speech classification with pseudo-labeling and task-specific knowledge.

7.2.1 Multi-View Pseudo-Labeling

Acoustic Path

We extract acoustic features from multiple audio encoders trained with different objectives to reduce the bias of relying on a single one, inspired by a previous study in the music field (Y. Li, Gui, Emmanouilidou, & Gamper, 2025). We use the following four audio encoders, resulting in four sets of embeddings for calculating the respective FAD scores:

- **VGGish:** convolutional embeddings (Hershey et al., 2017)
- **EnCodec:** low-rate audio codecs (Défossez, Copet, Synnaeve, & Adi, 2022)
- **wav2vec 2.0:** self-supervised acoustic embeddings (Baevski et al., 2020)
- **CLAP:** contrastive audio-text embeddings (Elizalde, Deshmukh, & Wang, 2024)

Since these models are pre-trained for different purposes, they are used to capture distinct acoustic characteristics, providing complementary embeddings that help reduce the bias in speech classification associated with relying on a single audio encoder.

Given the embeddings of labeled data \mathbf{X}^l and unlabeled data \mathbf{X}^u , the FAD score is calculated using multivariate Gaussians from two embedding sets $X^l(\mu_l, \Sigma_l)$ and $X^u(\mu_u, \Sigma_u)$ as follows:

$$F(X^l, X^u) = \|\mu_l - \mu_u\|^2 + tr(\Sigma_l + \Sigma_u - 2\sqrt{\Sigma_l \Sigma_u}) \quad (7.1)$$

where tr is the trace of a matrix.

Compared to traditional similarity metrics, such as cosine similarity and Euclidean distance, FAD is specifically designed for audio assessment, reflecting the perceptual similarity between two audio embedding distributions (Roblek, Kilgour, Sharifi, & Zuluaga, 2019). It has proven effective in distinguishing real and synthetic audio (Gui, Gamper, Braun, & Emmanouilidou, 2024), classifying audio with different emotions (Y. Li, Gui, et al., 2025), and measuring acoustic similarity between emotional music and speech (Sun, Zhao, Richmond, & Li, 2025). Therefore, we use FAD to measure the acoustic similarity between labeled and unlabeled data. The four FAD scores, calculated based on four sets of embeddings, are averaged to obtain the final FAD score. Finally, the unlabeled data with the smallest FAD score relative to the labeled class will be assigned that class label as the *acoustic pseudo-labels*. An example of this process for emotion pseudo-labeling is shown in Table 7.1. From the VGGish embedding, although the lowest individual FAD score corresponds to *Happy*, the lowest average score is for *Angry*. Thus, the unlabeled data is labeled as *Angry*. Note that the scores are not comparable across different audio encoders, as they are calculated based on different embeddings.

Table 7.1: An example of emotion pseudo-labeling using FAD. **Bold:** the smallest average FAD score, indicating its pseudo-label (i.e., Angry in this case).

	Angry	Happy	Neutral	Sad
VGGish	4.12	3.98	6.87	12.20
EnCodec	35.33	42.56	57.24	89.65
wav2vec 2.0	54.66	58.49	88.78	109.02
CLAP	45.46	182.65	141.75	230.39
Average	34.64	71.42	73.41	110.57

Linguistic Path

Previous speech classification tasks that used textual information typically relied on ground-truth text. However, in real-world applications, ASR is the only text source, and its transcriptions are usually noisy and contain errors, which can lead to incorrect classifications or pseudo-labels. We argue that *it is more challenging to prompt LLMs for classification tasks based on ASR transcriptions compared to human transcriptions due to the presence of word errors*.

To address this, we use the REVISE-REASON-RECOGNIZE (R3) prompting pipeline introduced in Section 5.2 to perform speech classification with ASR Error Correction (AEC) and reasoning on ASR transcriptions. As described, the R3 pipeline involves three steps: REVISE, where ASR errors are corrected based on N-best hypotheses; REASON, where the LLMs self-explain based on the corrected transcriptions and task-specific knowledge; and RECOGNIZE, where the label is identified.

For the ASR systems, we adopt the following ten models, the same as those used in Section 6.2, to generate diverse transcriptions and form 10-best ASR hypotheses:

- *wav2vec2-base-{100h,960h}*
- *wav2vec2-large-960h*
- *wav2vec2-large-960h-lv60-self*
- *HuBERT-large-ls960-ft*
- *WavLM-libri-clean-100h-base-plus*
- *Whisper-{tiny, base, small, large-v2}.en*

To perform AEC, we follow an AEC-specific Alpaca template (C.-H. H. Yang et al., 2023), which uses the “*You are an ASR error corrector*” instruction, guiding the LLMs to perform error correction. As LLMs have demonstrated their ability in both AEC and emotion recognition (C.-H. H. Yang et al., 2024), we expect that this capability can be extended to dementia detection from ASR transcriptions as well. The revised ASR transcriptions will be used for subsequent text feature extraction to train the bimodal classifier. For reasoning, we design task-specific knowledge that incorporates acoustics, linguistics, and psychology as in Figure 7.1.

We first apply Parameter-Efficient Fine-Tune (PEFT) on the following three LLMs with the LoRA adapter (E. J. Hu et al., 2021) using the labeled data with the R3 prompt:

- *Llama2-7b-chat-hf*
- *Llama2-13b-chat-hf*
- *Falcon-7b-instruct*

The objective of PEFT is to adapt a pretrained language model to either emotion or AD detection using a small number of task-specific parameters, while leveraging the knowledge from the pre-trained model. The optimization focuses on improving the model’s performance on the task at hand, typically by minimizing a task-specific loss function. By utilizing LoRA, PEFT achieves efficient fine-tuning by updating only a subset of the model’s parameters, which reduces computational cost and training time. The goal is to maximize task performance with minimal resource usage, especially when limited labeled data is available in this semi-supervised setting.

The learning rate, weight decay, and number of epochs are set to 1.0×10^{-4} , 1.0×10^{-5} , and 5, respectively, with AdamW optimizer used. The three fine-tuned LLMs are then prompted with the R3 prompt to predict class labels from ASR transcriptions of the unlabeled data. Finally, majority voting is applied to the predicted labels of the three LLMs to generate the *linguistic pseudo-labels*.

Data Selection

The acoustic and linguistic pseudo-labels generated from the two paths are combined to select the most confident data for semi-supervised training. Data with matching pseudo-labels from both paths are selected as high-confidence, while data with differing pseudo-labels are considered low-confidence. Together with the labeled data, the high-confidence data will be used to train the bimodal classifier in the first iteration, ensuring robust initial training.

7.2.2 Semi-Supervised Training

Bimodal Classifier

The bimodal classifier consists of pre-trained feature encoders: *HuBERT* (Hsu et al., 2021) (we use HuBERT here, a distinct model from those used in the audio path, to provide a new perspective, trying to avoid redundant information in the pseudo-labels generated by the classifier and the audio path) and *RoBERTa* (Y. Liu, 2019) that extract audio and text features, respectively, and a classification model that uses these features to generate a prediction label. For PEFT the encoders and training the classification model, the learning rate, weight decay, number of epochs, and batch size are set as 1.0×10^{-4} , 1.0×10^{-5} , 30, and 64, respectively. The AdamW optimizer is used. For the classification model, we examine the following four fusion methods:

- **Early fusion:** text and audio features are concatenated at the embedding level
- **Cross-attention fusion:** text and audio features are attended to each other via attention and then concatenated (Y. Li, Bell, & Lai, 2022)
- **Tensor fusion:** unimodal information and bimodal interactions are learned explicitly and then aggregated (Zadeh et al., 2017)
- **Modality-gated fusion:** primary modality is dynamically adjusted in each training step, as proposed in Section 6.3. As we only consider audio and text modalities, it aligns with that used in Section 6.4 for bimodal fusion

Iteration

After training the bimodal classifier, low-confidence data are predicted and labeled using the trained classifier¹. In most previous Semi-SL studies, model-labeled data are fully trusted and incorporated into the training set in the next iteration. However, as training progresses, mislabeled data (noise) may accumulate, leading to a cycle of erroneous learning (X. J. Zhu, 2005). To address this issue, we choose not to fully trust the model-labeled data. Instead,

1. Here, 'predicted' refers to the model generating a class prediction (i.e., assigning a probability distribution over emotion classes) for the unlabeled data, while 'labeled' refers to assigning the most probable class as a pseudo-label to that data point based on the model's prediction.

the pseudo-label generated by the bimodal classifier is compared with pseudo-labels from multi-view pseudo-labeling. If the model pseudo-label matches either the acoustic or linguistic pseudo-label, the data are added to the training set for the next iteration. Otherwise, they remain low-confidence and will be predicted in the next iteration.

In each iteration, we update the training set by adding model-labeled data and randomly removing 20% of the initial high-confidence data to avoid over-reliance on multi-view pseudo-labeling. The model is reinitialized in every iteration to reduce overfitting and bias. The maximum number of iterations is set to 40. However, if there is no performance improvement on the validation set for two consecutive iterations, the iteration will be terminated. The process is summarized in Algorithm 3.

Algorithm 3: Iteration Process

Input: H : Bimodal classifier; S_v : Validation set; D_c : Confident (high-confidence) data; D_u : Unconfident (low-confidence) data; L^a : Acoustic pseudo-labels; L^l : Linguistic pseudo-labels; L^h : Model pseudo-labels; I : Maximum number of iterations

```

1 for  $i = 1, \dots, I$  do
2   Train classifier  $H^i \leftarrow f(D_c^i)$ ;
3   Evaluate  $L_e \leftarrow H^i(S_v)$ ;
4   if performance on  $D_e$  does not improve then
5     | break;
6   end if
7   Generate pseudo-label  $L_h^i \leftarrow h^i(D_u)$ ;
8   if  $L_h^i$  equals  $L_a^i$  or  $L_l^i$  then
9     |  $L_u^i \leftarrow L_h^i$ ;
10  end if
11  Update  $D_c^i$  and  $D_u^i$ ;
12  Reinitialize classifier  $H^i$ ;
13 end for
14 end

```

7.3 Experimental Setup

Datasets

We use IEMOCAP for emotion recognition and ADReSSo (Luz, Haider, de la Fuente, Fromm, & MacWhinney, 2021) for dementia detection. For IEMOCAP, we focus on the Big Four emotion classes and exclude utterances with blank transcriptions, resulting in 5,500 utterances (1,103 angry, 1,615 happy+excited, 1,704 neutral, 1,078 sad). For ADReSSo, since there are no human labels in the test set to verify our approach, we use only the training set and focus on binary classes: Alzheimer’s Dementia (AD) and Cognitively Normal (CN). Due to the long duration of each audio file and the presence of interviewer’s speech, we segment all files using the official segmentation information, extracting participants’ speech, which results in

2,268 utterances (1,200 AD, 1,068 CN). Punctuation and extra whitespace are removed, and all text is converted to lowercase. The revised ASR transcriptions by R3 prompt yield word error rates of 11.48% and 30.25% for IEMOCAP and ADReSSo, respectively (the groundtruth transcription of ADReSSo are created by Whisper automatic transcription, followed by a native speaker’s correction, as there is no ground truth provided in the dataset).

Settings

For both tasks, we use an 80/10/10 split for training, validation, and testing, applying the ratio equally to each class to ensure a balanced distribution. Additionally, the ground-truth labeling rates for the training data are compared at 20%, 25%, and 30%, with the remaining data labeled using our method. All results are measured using Unweighted Accuracy (UA). Random seeds are kept consistent across all experiments. Other settings have been detailed in the previous section.

7.3.1 Results and Discussion

Four baselines are used for comparison:

- **Supervised_full**: the classification model is trained on the entire training data (i.e., the 80% split)
- **Supervised_limited**: the classification model is trained on the limited labeled data (20%, 25%, and 30%) without pseudo-labeling the unlabeled data
- **Decision merging**: two classification models are trained using audio and text, respectively, and their probability distribution are merged to select high-confidence data for the next iteration
- **Co-training**: two classification models are trained using audio and text, respectively. High-confidence data selected by each model are added to the training set for the other model in the next iteration (Blum & Mitchell, 1998; Z. Zhang et al., 2018)

Table 7.2: Performance comparison with *supervised_full* and *supervised_limited* (results in UA%). N : ground-truth labeling rate of training data. **BOLD**: best performance in each ground-truth labeling rate.

	Fusion	S_full	$N = 30%$		$N = 25%$		$N = 20%$	
			S_limited	Ours	S_limited	Ours	S_limited	Ours
Emotion recognition	<i>Early</i>	73.67	70.01	72.52 (+2.51)	69.76	72.04 (+2.28)	68.89	71.20 (+2.31)
	<i>C-attn</i>	74.02	71.20	73.79 (+2.59)	69.97	72.87 (+2.90)	69.03	71.85 (+2.82)
	<i>Tensor</i>	75.18	72.00	74.79 (+2.79)	70.42	73.25 (+2.83)	69.52	72.34 (+2.82)
	<i>M-gated</i>	75.53	72.03	75.10 (+3.07)	71.19	73.90 (+2.71)	70.88	73.61 (+2.73)
Dementia detection	<i>Early</i>	80.03	75.93	79.04 (+3.11)	73.55	78.75 (+5.20)	71.87	78.10 (+6.23)
	<i>C-attn</i>	80.32	76.77	79.99 (+3.22)	73.76	78.89 (+5.13)	71.18	78.25 (+7.07)
	<i>Tensor</i>	80.71	76.62	80.12 (+3.50)	74.82	79.48 (+4.66)	72.01	78.83 (+6.82)
	<i>M-gated</i>	81.23	77.14	80.87 (+3.73)	74.78	79.71 (+4.93)	72.14	79.11 (+6.97)

The comparison of our method with *supervised_full* and *supervised_limited* are shown in Table 7.2. It can be observed that

1. With only 30% labeled data, our method achieves performance competitive with the *supervised_full* baseline.
2. Our proposed method outperforms the *supervised_limited* baseline, which lacks multi-view pseudo-labeling and semi-supervised training to augment the labeled data, particularly in dementia detection.
3. The less ground-truth labeled data available, the more effective our method is in dementia detection, likely because binary classes are easier for pseudo-labeling.
4. When ground-truth labels are limited, classification performance of the *supervised_limited* baseline for both tasks drops largely compared to the *supervised_full* training. Our proposed method, however, mitigates this drop by more than 2% in emotion recognition and 3%-7% in dementia detection.
5. Modality-gated fusion performs best among the four fusion methods. This is reasonable as it dynamically selects the primary modality contributing most to the classification tasks, thereby reducing the impact of ASR errors in the text modality.

As bimodal fusion does not apply to the *decision merging* and *co-training* baseline settings, we select our best-performing results for comparison. Note that the principle of *decision merging* is the same as that of a fusion method: late fusion (or decision-level fusion) (Snoek, Worring, & Smeulders, 2005). The difference is that we further use the merged probability to determine the high-confidence data for iteration, whereas late fusion directly outputs results based on the highest probability. Here, we refer to it as *decision merging* to avoid confusion with late fusion, as we have used fusion techniques in the previous experiment.

For both *decision merging* and *co-training*, we set the threshold at 0.5 for emotion recognition and 0.7 for dementia detection. For example, the fourth emotion class will be selected as the pseudo-label for an unlabeled data with a probability distribution of {0.1, 0.1, 0.3, 0.5}, and the unlabeled data will be added to the training set as high-confidence data. In general, a lower threshold is more suitable for multi-class tasks, while a higher threshold is preferred in binary tasks with higher risk of misclassification. Therefore, using 0.5 for emotion recognition and 0.7 for dementia detection is reasonable and aligns with confidence-based rejection methods in previous studies (Chow, 1970; Geifman & El-Yaniv, 2017; H. Jiang, Kim, Guan, & Gupta, 2018), where a reasonably high threshold is typically selected to balance accuracy and rejection rate.

From Table 7.3, we can observe that:

Table 7.3: Performance comparison with *decision merging* and *co-training* (results in UA%). N : ground-truth labeling rate of training data. **BOLD**: best performance in each rate.

		Emotion recognition	Dementia detection
$N = 30\%$	<i>Decision merging</i>	67.60	74.55
	<i>Co-training</i>	69.99	74.73
	Ours	75.10	80.87
$N = 25\%$	<i>Decision merging</i>	66.34	72.70
	<i>Co-training</i>	69.04	72.87
	Ours	73.90	79.71
$N = 20\%$	<i>Decision merging</i>	66.12	71.07
	<i>Co-training</i>	69.00	72.10
	Ours	73.61	79.11

1. Our method largely outperforms the two baselines, likely for two reasons: It generates higher-confidence unlabeled data (since the pseudo labels come not only from the classifier but also from the audio and text pseudo labels, enhancing reliability), which iteratively trains the classifier to improve performance; second, our bimodal classifier takes both modalities as input during training. In contrast, although the two baselines also consider two modalities, their classifiers are separate, each relying on a single modality and ignoring the interrelatedness between them.
2. Although both *decision merging* and *co-training* consist of two classifiers that output respective labels, the latter performs better than the former, especially in emotion recognition. This result is plausible since *decision merging* could potentially weaken the better prediction if the other probability is greatly incorrect. On the contrary, by incorporating high-confidence data judged by the other view into training set, *co-training* enables the two models to learn from each other indirectly. Its relatively lower effectiveness in dementia detection is likely due to the low-quality audio, which makes one of the models less powerful.

Effect of Multi-View Pseudo-Labeling Components

We explore the contributions of each audio encoder and LLM in multi-view pseudo-labeling by keeping either the acoustic or linguistic path unchanged (i.e., full path) while adding an individual encoder from the other path. For brevity, Table 7.4 presents the results of early fusion with 30% labeled data omitting the other conditions. The results show that **1)** The acoustic path contributes more than the linguistic path. **2)** *CLAP* and *Falcon* perform best among the acoustic and linguistic encoders, respectively.

Table 7.4: Contribution of each component of the multi-view pseudo-labeling under the scenario of 30% labeled data (results in UA%).

	Emotion recognition	Dementia detection
Two paths	72.52	79.04
Full linguistic path		
+ <i>VGGish</i>	68.21	72.04
+ <i>EnCodec</i>	65.75	70.76
+ <i>wav2vec 2.0</i>	68.87	72.00
+ <i>CLAP</i>	69.41	72.62
Full acoustic path		
+ <i>Llama2-7b</i>	70.53	76.40
+ <i>Llama2-13b</i>	71.31	77.91
+ <i>Falcon</i>	72.00	78.88

7.4 Chapter Summary

In this chapter, we proposed a novel semi-supervised learning framework that introduces a multi-view pseudo-labeling method leveraging both acoustic and linguistic characteristics. This method utilized Fréchet audio distance and large language models to select the most reliable unlabeled data for augmenting the training set. Multiple fusion techniques have been compared to utilize multi-view knowledge for further enhancement of the framework. We evaluated our method on emotion recognition and dementia detection tasks, demonstrating that it outperforms fully-supervised, limited-supervised, and two SSL baselines. Our method achieved competitive performance compared to fully supervised learning while using less than 30% of human-labeled data.

Compared to existing literature on semi-supervised training for speech classification tasks, we introduced a novel pseudo-labeling framework that leverages audio distance and outputs from LLMs. This approach improves the reliability of pseudo labels, as their confidence is informed by sources beyond the emotion classifier itself. As a result, the issue of error accumulation common in traditional semi-supervised learning algorithms is largely mitigated.

Several limitations remain, which we plan to address in future work: 1) exploring a wider range of audio encoders and LLMs to identify more reliable combinations; 2) refining the confidence threshold to better balance accuracy and the amount of high-confidence data for iteration; and 3) extending the method to additional speech classification tasks, especially those involving low-quality audio.

Conclusions and Outlook

8.1 Review of Thesis Achievements

Speech emotion is a natural, complex, and central component of human verbal communication. Analyzing and understanding the emotion in human speech is a crucial function for Artificial General Intelligence (AGI). The implementation of SER in real-world applications has gained traction across various industries (Y. Li, 2022, 2023b). This thesis aims to accelerate this by integrating ASR functions.

With the evolution of speech technologies, the way of using speech models in SER has progressed with time. The premise of utilizing state-of-the-art speech foundation models in SER is understanding whether and how they function on emotional speech.

To address the first research question outlined in Chapter 1, “*What is the interplay between ASR models and SER?*”, this thesis examines the behavior of a state-of-the-art speech foundation model, wav2vec 2.0, on emotional speech through a series of comprehensive experiments presented in Section 3.2. The results indicate that the representations generated by wav2vec 2.0, which are optimized for ASR tasks, often lack critical paralinguistic cues and exhibit distinct hierarchical structures and emotion biases. In Section 4.2, we further investigate how word distribution within a corpus influences ASR performance, uncovering notable emotion-specific variations. Factors such as PoS tag frequency, affective ratings, and utterance length were shown to greatly affect ASR accuracy. Certain word categories may be more easily recognized by one model but pose challenges for another. Additionally, Section 4.3 presents an evaluation of SER systems using ASR-generated transcripts with varying WERs. We also examine common fusion strategies to understand how ASR quality impacts SER outcomes. To this end, we achieved the two objectives listed under this research question: 1) Analyzing how ASR systems process emotionally expressive speech and quantify their limitations; 2) Examining the impact of ASR errors on downstream SER models, particularly in real-world, spontaneous speech settings.

To address the second research question “*How to improve ASR transcription quality of small-scale emotional speech for SER?*”, Chapter 5 introduces both LLM-based and S2S-based approaches tailored to different ASR hypothesis scenarios, namely, N-best and 1-best outputs. Section 5.2 presents a method based on LLMs, which refines N-best ASR hypotheses through emotion-aware prompts and ASR error correction, ultimately enhancing emotion recognition from the improved transcriptions. In Section 5.3, a S2S model is proposed that integrates discrete acoustic units with textual word units to perform cross-modal ASR error correction, outperforming the LLM-based method when applied to 1-best ASR outputs. The transcriptions enhanced by both methods have been utilized to address the third question, demonstrating their effectiveness in delivering higher-quality textual input for SER. To this end, the two objectives listed under this research question have been realized: 1) Adapting existing ASR error correction methods to low-resource emotional speech domains, with a focus on learning to correct emotion-specific transcription errors; 2) Improving SER performance by utilizing corrected ASR transcriptions in both text-only and speech-text fusion approaches.

In addressing the third research question “*Can we develop robust acoustic-lexical fusion frameworks for ASR-enhanced SER to mitigate error propagation?*”, Section 3.3 presents a joint ASR-SER training framework that hierarchically incorporates eGeMAPS paralinguistic features, intermediate representations from wav2vec 2.0, and ASR-generated transcriptions into the SER process. This design leverages the strengths of wav2vec 2.0 while addressing its limitations, resulting in performance comparable to that achieved using ground-truth transcripts. Furthermore, Section 6.2 examines a key challenge in multimodal fusion: cross-modal incongruity. To address this, Section 6.3 introduces an incongruity-aware fusion strategy that helps the model discern the dominant affective tendency across modalities. Section 6.4 then presents an ASR error-robust framework combining ASR error correction with modality-gated fusion. This approach outperforms baseline models in our experiments and shows promise for text-only and bimodal fusion SER tasks where ASR transcriptions provide the textual input. Collectively, these contributions fulfill the two objectives: 1) Designing robust ASR-enhanced SER frameworks that effectively integrate speech and text cues while mitigating the impact of ASR errors; 2) Investigating novel fusion techniques, such as error-aware fusion and adaptive attention mechanisms to enhance SER robustness in real-life scenarios.

In response to the final research question “*Can effort-saving strategies be developed for scalable SER to reduce annotation costs?*”, Chapter 7 introduces an innovative semi-supervised learning framework based on a multi-view pseudo-labeling approach that incorporates both acoustic and linguistic cues. This framework employs Fréchet audio distance and LLMs to identify and select the most trustworthy unlabeled samples for training augmentation. Several fusion strategies are explored to integrate multi-view knowledge and further enhance model performance. Experiments on emotion recognition and dementia detection tasks demonstrate that the proposed approach surpasses baseline methods. Distinct from prior semi-supervised

efforts in speech classification, our pseudo-labeling strategy combines audio similarity metrics and LLM outputs, yielding more reliable pseudo labels by incorporating external indicators of confidence beyond the emotion classifier itself. Consequently, the framework fulfills three objectives: 1) Developing novel SemiSL strategies for SER, incorporating both confidence-based and reliability-aware sample selection; 2) Exploring co-training methodologies that exploit multi-view learning across various feature encoders; 3) Reducing dependence on large-scale annotated emotion datasets and leverage ASR transcription, making SER models more cost-efficient and effort-saving.

8.2 Toward Reliable and Applicable SER

Despite the achievements presented in this thesis, several limitations remain and warrant further investigation.

First, as discussed in Section 3.2, speech foundation models such as wav2vec 2.0, although highly effective in ASR and other speech-related tasks, exhibit limited sensitivity to paralinguistic cues. These models are primarily optimized for phonetic accuracy and thus tend to abstract away emotion-relevant acoustic characteristics. This inherent bias poses a challenge for their direct application in emotion-sensitive domains such as mental health monitoring, customer service, and social robotics, where nuanced emotional perception is critical.

Second, while the ASR error correction strategies proposed in Chapter 5 improve transcription quality, their effectiveness remains highly dependent on the domain similarity between training data and deployment scenarios. Additionally, many emotionally salient acoustic events, such as vocal bursts, laughter, or sighs, are typically excluded from ASR transcriptions and cannot be adequately recovered through text-based correction methods. These omissions limit the emotional fidelity of the textual modality and may introduce blind spots in SER systems.

Third, although the fusion strategies developed in this thesis successfully integrate speech and text modalities, emotional signals across modalities are not always congruent, particularly in instances of sarcasm, irony, or emotionally masked expressions. Such cases demand a deeper level of contextual and pragmatic understanding. For instance, while Chapter 6 explores incongruity-aware fusion, the interpretability of the model remains limited. It is still unclear whether the model genuinely understands cross-modal incongruity or merely exploits superficial statistical patterns.

Finally, scalability poses an ongoing challenge. Although the LLM-based techniques introduced in Chapter 5 and the semi-supervised learning framework in Chapter 7 offer promising avenues for reducing manual annotation costs, their success relies on the capabilities of the underlying LLMs and the reliability of the emotion labels they produce. The mechanisms

by which LLMs infer emotions from textual input remain questionable, and their failures are often difficult to diagnose. In high-stakes applications such as healthcare or security, even small errors can have serious consequences, highlighting the urgent need for more robust confidence estimation and model interpretability.

We are currently exploring several of these challenges in ongoing collaborations, with the aim of advancing toward more reliable and applicable SER systems. Future research should prioritize the following directions:

- Developing emotion-centric pretraining objectives for speech foundation models that better preserve affective and paralinguistic information.
- Advancing multilingual and cross-cultural emotion modeling, including robust handling of code-switching and culturally specific affective expressions.
- Designing context-aware fusion architectures capable of dynamically assessing and adapting to the reliability of different modalities.
- Implementing uncertainty quantification and calibration techniques to improve robustness and trustworthiness in SER and relevant applications.
- Conducting ethical and fairness assessments to ensure that SER systems do not propagate bias or misrepresent underrepresented demographic groups.

By addressing these limitations, we can pave the way for SER technologies that are not only technically robust but also socially responsible and truly applicable in real-world contexts.

Research Outputs Related to This Thesis

In addition to the papers directly stemming from this thesis as described in Section 1.3, there are also several papers related to this work. Below describes how they contribute to the topics relevant to this thesis.

1. Semi-Supervised Learning for Multimodal Speech and Emotion Recognition (Y. Li, 2021b)

This paper is the first submission of the PhD work, intended for the Doctoral Consortium at ICMI 2021. It outlines the planned direction of the thesis, proposing the ideas of hierarchical fusion and multi-view semi-supervised training, which lay the foundation for a significant portion of the thesis.

2. Alzheimer’s Dementia Detection through Spontaneous Dialogue with Proactive Robotic Listeners (Y. Li, Lai, Lala, Inoue, & Kawahara, 2022)

This paper marks our initial exploration into speech-based detection of Alzheimer’s disease, accepted to IEEE HRI 2022 late-breaking session. We proposed a method that integrates both speech and text information by utilizing the feature of “focus words.” The concept of focus words later inspired our use of trigger words in LLMs for emotion recognition and Alzheimer’s dementia detection in the thesis work.

3. Robotic Speech Synthesis: Perspectives on Interactions, Scenarios, and Ethics (Y. Li & Lai, 2022b)

This paper discusses the role of prosody in emotional speech, accepted to IEEE HRI Robo-identity 2 Workshop¹, which provides a foundation for our subsequent exploration of whether embeddings from speech foundation models can encode emotion-related prosodic characteristics in the thesis work. Additionally, the paper addresses ethical issues surrounding synthesized speech, which inspired the special session “Responsible Speech Foundation Models” that we later organized at Interspeech (Y. Li, Williams, et al., 2023).

1. <https://sites.google.com/view/robo-identity2>

4. A Cross-Domain Approach for Continuous Impression Recognition from Dyadic Audio-Visual-Physio Signals (Y. Li & Lai, 2022a)

This paper is our solution submitted to IUI 2022-Dyadic IMPRESSION Recognition Challenge². We integrated multimodal signals from both the speaker and the listener, including speech, eye gaze, facial expressions, and various physiological signals. We proposed two novel loss functions, Knowledge Distillation Loss and Similarity Enhancement Loss, to recognize the listener's impression when receiving information from the speaker, outperforming the challenge baseline. It serves as an extension of the multimodal fusion work presented in this thesis.

5. Multimodal Dyadic Impression Recognition via Listener Adaptive Cross-Domain Fusion (Y. Li, Bell, & Lai, 2023a)

This paper extends the above work by introducing a novel listener adaptation module, which measures and leverages the causal relationship between the speaker's and listener's signals as weights for multimodal fusion. This method further improved the performance over the prior approach and has been accepted to IEEE ICASSP 2023. The weight-adaptation multimodal fusion partly inspired our modality-gated fusion in the thesis work.

6. Transfer Learning for Personality Perception via Speech Emotion Recognition (Y. Li, Bell, & Lai, 2023b)

This paper explores transferring knowledge from an emotion recognition model based on wav2vec 2.0 to the task of personality detection, aiming to address the challenge of limited data in speech-based tasks. This work extends our thesis research on the wav2vec 2.0 model, further validating its applicability to emotion-related tasks. It has been accepted to Interspeech 2023.

7. I Know Your Feelings Before You Do: Predicting Future Affective Reactions in Human-Computer Dialogue (Y. Li, Inoue, et al., 2023)

In this paper, we proposed an emotion prediction framework that operates prior to emotion recognition in dialogue systems. By leveraging the speaker's emotional and linguistic features (such as dialog acts), the system can predict the listener's forthcoming emotional state. This approach contributes to smoother emotional dialogues and reduces computational resource usage. We analyzed real human emotional conversations to validate the feasibility of this approach. This work represents an attempt to apply our PhD research to real-world scenarios, and has accepted to CHI 2023 late-breaking work.

8. Empowering Dialogue Systems with Affective and Adaptive Interaction: Integrating Social Intelligence (Y. Li & Lai, 2023)

2. <https://competitions.codalab.org/competitions/34190>

This paper builds upon the foundational work of this thesis and explores how to incorporate emotional intelligence and social intelligence into dialogue systems, enabling SER systems to be used in truly applicable ways. We proposed that dialogue systems should either simulate the user's emotions or maintain their own style, depending on the scenarios. This paper has been accepted to ACII 2023 Social and Affective Intelligence Workshop and won the outstanding paper award.

9. Enhancing Speech Emotion Recognition for Real-World Applications via ASR Integration (Y. Li, 2023a)

This paper is another submission to a Doctoral Consortium (at ACII 2023), outlining ongoing work and emphasizing the importance of integrating ASR into SER. It is a revision of our initial doctoral consortium paper presented at ICMI 2021 and lays the foundation for our later work on semi-supervised training.

10. Layer-Wise Analysis of Self-Supervised Acoustic Word Embeddings: A Study on Speech Emotion Recognition (Saliba et al., 2024)

This work is a collaborative study between me and the master's student Alexandra Saliba that I supervised in 2023. The work builds on our exploration of wav2vec 2.0 in the thesis. Instead of using continuous representations, we opted for Discrete Speech Units (DSUs), considering two key advantages they offer: First, a significant reduction in computational resources. Unlike the continuous representations with dimensions of (frame number, embedding), DSUs only have dimensions of (1, embedding). Second, DSUs have been shown to capture certain lexical information. We conducted a layer-wise analysis and applied DSUs to SER, which confirmed their effectiveness. It has been accepted to ICASSP 2024 Self-supervision in Audio, Speech and Beyond Workshop.

11. Beyond Voice Assistants: Exploring Advantages and Risks of an In-Car Social Robot in Real Driving Scenarios (Y. Li, Urquhart, et al., 2024)

This paper represents another exploration of practical applications for speech emotion dialogue systems. We utilized the in-car voice assistant of a mass-produced vehicle to investigate how the WER of ASR impacts user perception, as well as how the speech personality and emotional dimensions of the in-car assistant influence the driving experience. This work was partially carried out prior to my PhD, during my time at Honda Innovation Lab. It connects my earlier industry research with my current academic pursuits, expanding our thesis work on ASR and SER into the domain of human-computer interaction.

12. Chance-Constrained Abnormal Data Cleaning for Robust Classification With Noisy Labels (X. Shen, Luo, Li, Ouyang, & Wu, 2024)

This is a collaborative work on the reliability of emotion labels, exploring data-cleaning methods to achieve robust classification despite noisy labels. Rather than using all the labels from a corpus, our approach focuses on filtering out abnormal data to improve classification tasks. We tested our method on IEMOCAP and enhanced the performance of supervised classification methods. It partly inspires the multi-view pseudo-labeling approach in this thesis to improve data reliability. This paper has been accepted to IEEE Transactions on Emerging Topics in Computational Intelligence.

13. 1st Place Solution to Odyssey Emotion Recognition Challenge Task1: Tackling Class Imbalance Problem (M. Chen et al., 2024)

This paper is our solution submitted to Odyssey Emotion Recognition Challenge³. We introduced an ensemble system that includes seven multi-modal models, which were trained independently with different loss functions and class weights. Specifically, the cross entropy loss and the focal loss were used. The system showed the best performance among a total of 68 submissions to the challenge, in all metrics. The S2S-based ASR error correction function and several fusion approaches in this thesis were used in this paper.

14. Rethinking Emotion Bias in Music via Frechet Audio Distance (Y. Li, Gui, et al., 2025)

This paper presents the work I conducted during my internship at Microsoft Research, where I extended part of my thesis research from speech to the domain of music. In this project, we established a benchmark for music emotion recognition by evaluating multiple music datasets and various audio encoders. Additionally, we improved two music generation models by integrating our proposed emotion generation module, which resulted in the production of more emotionally realistic music. We explored the biases introduced by different metrics, recognition tasks, and models, and proposed using Frechet Audio Distance as a solution. This method is applicable not only in the domain of music emotion but also in speech emotion and dementia analysis later studied in this thesis work. This paper has been accepted to IEEE ICME 2025.

15. Exploring Acoustic Similarity in Emotional Speech and Music via Self-Supervised Representations (Sun et al., 2025)

This work is a collaborative study between me and the master's student Yujia Sun that I supervised in 2024. It is an extension of the wav2vec 2.0 exploration in this thesis. We compared two speech foundation models (wav2vec 2.0 and HuBERT) and one music foundation model (MERT) in terms of their performance on emotional speech and emotional music. We found that while speech and music SSL models do capture shared acoustic features, their behaviors can vary depending on different emotions. This paper has been accepted to IEEE ICASSP 2025.

3. <https://www.odyssey2024.org/emotion-recognition-challenge>

16. Cross-Lingual Speech Emotion Recognition: Humans vs. Self-Supervised Models (Z. Han et al., 2025)

This work is a collaborative study between me and the master's student Zhichen Han that I supervised in 2024. It extends this thesis research by exploring the comparison between self-supervised speech models and human performance in cross-lingual SER. We used Mandarin Chinese, English, German, and Tianjin dialect (a Chinese regional dialect), along with three speech foundation models, to investigate how the same model performs in SER across different languages, and how it can be fine-tuned on the target language. The study found that a model trained on language A can, with simple fine-tuning on language B, achieve performance comparable to that of humans. We also discovered that speech foundation models are capable of identifying emotional segments within speech in cross-lingual tasks. This paper has also been accepted to IEEE ICASSP 2025.

17. Segmentation-Variant Codebooks for Preservation of Paralinguistic and Prosodic Information (Sanders, Li, Richmond, & King, 2025)

This work is a collaborative study between me and another PhD student Nicholas Sanders in 2024. It extends Chapter 5 by exploring the effectiveness of discrete speech units of preserving prosodic and paralinguistic information in emotion recognition. We proposed novel segmentation-variant codebooks, which quantize speech at distinct linguistic units (frame, phone, word, utterance), factorizing it into multiple streams of segment-specific discrete features. The results show that the segmentation-variant codebooks are significantly more effective at preserving prosodic and paralinguistic information across probing tasks, including emotion recognition. This paper has been accepted to Interspeech 2025.

18. Leveraging Cascaded Binary Classification and Multimodal Fusion for Dementia Detection through Spontaneous Speech (Y.-L. Liu et al., 2025)

This paper is a collaborative study between me and the researchers at University of Science and Technology of China. It extends the multimodal fusion techniques discussed in Chapter 4 to Alzheimer's disease detection, and we submitted to the PROCESS Challenge 2025. For the three-class classification task (healthy control, mild cognitive impairment, and dementia), we propose a cascaded binary classification framework that fine-tunes pre-trained language models and incorporates pause encoding to better capture disfluencies. Experimental results on the test set outperform the baselines provided by the challenge in both tasks, demonstrating the robustness and effectiveness of our approach. This paper has also been accepted to Interspeech 2025.

19. Exploring Gender Bias in Alzheimer's Disease Detection: Insights from Mandarin and Greek Speech Perception (L. He et al., 2025)

This paper is a collaborative study between me and the researchers at University of Science and Technology of China. It extends the cross-lingual Alzheimer’s research in Chapter 7. In a perception experiment involving 16 Chinese listeners evaluating both Chinese and Greek speech, we identified that male speech was more frequently identified as AD, with this bias being particularly pronounced in Chinese speech. Acoustic analysis showed that shimmer values in male speech were significantly associated with AD perception, while speech portion exhibited a significant negative correlation with AD identification. This paper has been accepted to National Conference on Man-Machine Speech Communication 2025.

20. HRAI 2025: The 1st Workshop on Holistic and Responsible Affective Intelligence (Y. Li, Kollias, et al., 2025)

This paper summarizes the workshop I led at ICMI 2025⁴. It builds upon my entire thesis work, aiming to foster discussion and contributions toward holistic and responsible affective intelligence through advanced foundation models and machine learning techniques. The paper has been accepted to ICMI 2025.

21. Large Language Model Based Generative Error Correction: A Challenge and Baselines for Speech Recognition, Speaker Tagging, and Emotion Recognition (C.-H. H. Yang et al., 2024)

This paper presents the challenge summary and baseline systems for our hosted competition at IEEE SLT 2024⁵. It is partially derived from the ASR error correction research discussed in Chapter 5. We provide baseline methods using LLMs for ASR error correction in speech recognition, speaker tagging, and emotion recognition tasks.

4. <https://icmi.acm.org/2025/workshops/>

5. <https://2024.ieeeslt.org/challenges/#1715507565729-916ec1d3-b60d>

Bibliography

- Acosta, J. C., & Ward, N. G. (2011). Achieving rapport with turn-by-turn, user-responsive emotional coloring. *Speech Communication*, 53(9-10), 1137–1148.
- Amin, M. M., & Schuller, B. W. (2024). On prompt sensitivity of chatgpt in affective computing. *arXiv preprint arXiv:2403.14006*.
- Amiriparian, S., Sokolov, A., Aslan, I., Christ, L., Gerczuk, M., Hübner, T., ... others (2021). On the impact of word error rate on acoustic-linguistic speech emotion recognition: An update for the deep learning era. *arXiv preprint arXiv:2104.10121*.
- Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., ... others (2016). Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning* (pp. 173–182).
- Ancilin, J., & Milton, A. (2021). Improved speech emotion recognition with mel frequency magnitude coefficient. *Applied Acoustics*, 179, 108046.
- Ando, A., Masumura, R., Sato, H., Moriya, T., Ashihara, T., Ijima, Y., & Toda, T. (2021). Speech emotion recognition based on listener adaptive models. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6274–6278).
- Ardila, R., Branson, M., Davis, K., Henretty, M., Kohler, M., Meyer, J., ... Weber, G. (2020). Common voice: A massively-multilingual speech corpus. *Proceedings of the Twelfth Language Resources and Evaluation Conference (LREC)*.
- Attardo, S. (1994). *Linguistic theories of humor*. Berlin: Mouton de Gruyter.
- Ba, J. L., Kiros, J. R., & Hinton, G. E. (2016). Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Baevski, A., Zhou, Y., Mohamed, A., & Auli, M. (2020). wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33, 12449–12460.
- Baken, R. J., & Orlikoff, R. F. (2000). *Clinical measurement of speech and voice* (2nd ed.). Cengage Learning.
- Bandhakavi, A., Wiratunga, N., Massie, S., & Cooper, P. (2016). Lexicon based feature extraction for emotion text classification. In *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis* (pp. 89–95).

- Barsties, B., & Maryn, Y. (2013). The alpha ratio as a voice parameter. *Journal of Voice*, 27(4), 401–408.
- Bird, S. (2006). NLTK: the natural language toolkit. In *Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions* (pp. 69–72).
- Blum, A., & Mitchell, T. (1998). Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on computational learning theory* (pp. 92–100).
- Boersma, P. (1993). Accurate short-term analysis of the fundamental frequency and the harmonics-to-noise ratio of a sampled sound. *Proceedings of the Institute of Phonetic Sciences*, 17, 97–110.
- Boigne, J., Liyanage, B., & Östrem, T. (2020). Recognizing more emotions with less data using self-supervised transfer learning. *arXiv preprint arXiv:2011.05585*.
- Busso, C., Bulut, M., Lee, C.-C., Kazemzadeh, A., Mower, E., Kim, S., ... Narayanan, S. S. (2008). IEMOCAP: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42, 335–359.
- Busso, C., Lee, S., & Narayanan, S. S. (2007). Using neutral speech models for emotional speech analysis. In *Proceedings of Interspeech 2007*.
- Cai, X., Yuan, J., Zheng, R., Huang, L., & Church, K. (2021). Speech emotion recognition with multi-task learning. In *Proceedings of Interspeech 2021*.
- Cambria, E., Livingstone, A., & Hussain, A. (2012). The hourglass of emotions. In *Cognitive behavioural systems: Cost 2102 international training school, dresden, germany, february 21-26, 2011, revised selected papers* (pp. 144–157).
- Cauci, G. M., & Kreuz, R. J. (2012). Social and paralinguistic cues to sarcasm. *Humor*, 25(1), 1–22.
- Cettolo, M., Girardi, C., & Federico, M. (2012). WIT3: Web inventory of transcribed and translated talks. In *Proceedings of the 16th annual conference of the european association for machine translation* (pp. 261–268). Trento, Italy: European Association for Machine Translation.
- Chan, W., Jaitly, N., Le, Q., & Vinyals, O. (2016). Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 4960–4964).
- Chatterjee, A., et al. (2020). Semeval-2019 task 3: Emocontext contextual emotion detection in text. *arXiv preprint arXiv:1902.01091*.

- Chauhan, D. S., Akhtar, M. S., Ekbal, A., & Bhattacharyya, P. (2019). Context-aware interactive attention for multi-modal sentiment and emotion analysis. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 5647–5657).
- CHEN, C., Li, R., Hu, Y., Siniscalchi, S. M., Chen, P.-Y., Chng, E., & Yang, C.-H. H. (n.d.). It's Never Too Late: Fusing Acoustic Information into Large Language Models for Automatic Speech Recognition. In *The Twelfth International Conference on Learning Representations*.
- Chen, L.-W., & Rudnicky, A. (2023). Exploring wav2vec 2.0 fine tuning for improved speech emotion recognition. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Chen, M., Zhang, H., Li, Y., Luo, J., Wu, W., Ma, Z., . . . others (2024). 1st place solution to odyssey emotion recognition challenge task1: Tackling class imbalance problem. In *Proceedings of odyssey 2024* (pp. 260–265).
- Chen, S., Wang, C., Chen, Z., Wu, Y., Liu, S., Chen, Z., . . . others (2022). Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6), 1505–1518.
- Chen, W., Xing, X., Xu, X., Pang, J., & Du, L. (2022). SpeechFormer: A Hierarchical Efficient Framework Incorporating the Characteristics of Speech. *arXiv preprint arXiv:2203.03812*.
- Chou, H.-C. (2024). A Tiny Whisper-SER: Unifying Automatic Speech Recognition and Multi-label Speech Emotion Recognition Tasks. In *2024 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)* (pp. 1–6).
- Chou, H.-C., Wu, H., Goncalves, L., Leem, S.-G., Salman, A., Busso, C., . . . Lee, C.-C. (2024). Embracing Ambiguity And Subjectivity Using The All-Inclusive Aggregation Rule For Evaluating Multi-Label Speech Emotion Recognition Systems. In *2024 IEEE Spoken Language Technology Workshop (SLT)* (pp. 502–509).
- Chow, C. (1970). On optimum recognition error and reject tradeoff. *IEEE Transactions on Information Theory*, 16(1), 41–46.
- Chuang, Z.-J., & Wu, C.-H. (2004). Multi-modal emotion recognition from speech and text. In *International Journal of Computational Linguistics & Chinese Language Processing, Volume 9, Number 2, August 2004: Special Issue on New Trends of Speech and Language Processing* (pp. 45–62).

- Chung, Y.-A., Hsu, W.-N., Tang, H., & Glass, J. (2019). An unsupervised autoregressive model for speech representation learning. In *Proceedings of Interspeech 2019* (pp. 146–150).
- Chung, Y.-A., Tang, H., & Glass, J. (2020). Vector-quantized autoregressive predictive coding. In *Proceedings of Interspeech 2020* (pp. 3760–3764).
- Church, K. W. (2017). Word2vec. *Natural Language Engineering*, 23(1), 155–162.
- Coutinho, E., & Dibben, N. (2013). Psychoacoustic cues to emotion in speech prosody and music. *Cognition & emotion*, 27(4), 658–684.
- Cowie, R., & Cornelius, R. R. (2003). Describing the emotional states that are expressed in speech. *Speech communication*, 40(1-2), 5–32.
- Cowie, R., Douglas-Cowie, E., Tsapatsoulis, N., Votsis, G., Kollias, S., Fellenz, W., & Taylor, J. G. (2001). Emotion recognition in human-computer interaction. *IEEE Signal processing magazine*, 18(1), 32–80.
- Cummins, N., Baird, A., & Schuller, B. W. (2018). Speech analysis for health: Current state-of-the-art and the increasing impact of deep learning. *Methods*, 151, 41–54.
- Daneshfar, F., Kabudian, S. J., & Neekabadi, A. (2020). Speech emotion recognition using hybrid spectral-prosodic features of speech signal/glottal waveform, metaheuristic-based dimensionality reduction, and gaussian elliptical basis function network classifier. *Applied Acoustics*, 166, 107360.
- Défossez, A., Copet, J., Synnaeve, G., & Adi, Y. (2022). High fidelity neural audio compression. *Transactions on Machine Learning Research*.
- Degottex, G., Kane, J., Drugman, T., Raitio, T., & Scherer, S. (2014). Covarep—a collaborative voice analysis repository for speech technologies. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 960–964).
- Deng, L., Wang, Y.-Y., Wang, K., Acero, A., Hon, H.-W., Droppo, J. G., . . . Huang, X. (2004). Speech and language processing for multimodal human-computer interaction. *Real World Speech Processing*, 87–113.
- Desai, P., Chakraborty, T., & Akhtar, M. S. (2022). Nice perfume. How long did you marinate in it? Multimodal Sarcasm Explanation. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 36, pp. 10563–10571).
- Devillers, L., Tahon, M., Sehili, M. A., & Delaborde, A. (2015). Inference of human beings' emotional states from speech in human–robot interactions. *International Journal of Social Robotics*, 7, 451–463.

- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT 2019*, 4171–4186.
- Diatlova, D., Udalov, A., Shutov, V., & Spirin, E. (2024). Adapting WavLM for Speech Emotion Recognition. In *Proceedings of odyssey 2024* (pp. 303–308).
- d'Sa, A. G., Illina, I., Fohr, D., Klakow, D., & Ruiter, D. (2020). Label propagation-based semi-supervised learning for hate speech classification. In *Insights from Negative Results Workshop, EMNLP 2020*.
- Du, J., Pu, S., Dong, Q., Jin, C., Qi, X., Gu, D., ... Zhou, H. (2022). Cross-modal ASR post-processing system for error correction and utterance rejection. *arXiv preprint arXiv:2201.03313*.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., ... others (2024). The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Dutta, S., Jain, S., Maheshwari, A., Pal, S., Ramakrishnan, G., & Jyothi, P. (2022). Error correction in ASR using sequence-to-sequence models. *arXiv preprint arXiv:2202.01157*.
- Ekman, P., Dalglish, T., & Power, M. (1999). Basic emotions. *San Francisco, USA*.
- El Ayadi, M., Kamel, M. S., & Karray, F. (2011). Survey on speech emotion recognition: Features, classification schemes, and databases. *Pattern recognition*, 44(3), 572–587.
- Elizalde, B., Deshmukh, S., & Wang, H. (2024). Natural language supervision for general-purpose audio representations. In *2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 336–340).
- Elsayed, N., ElSayed, Z., Asadizanjani, N., Ozer, M., Abdelgawad, A., & Bayoumi, M. (2022). Speech emotion recognition using supervised deep recurrent system for mental health monitoring. In *2022 IEEE 8th World Forum on Internet of Things (WF-IoT)* (pp. 1–6).
- Eyben, F., Scherer, K. R., Schuller, B. W., Sundberg, J., André, E., Busso, C., ... others (2015). The Geneva minimalistic acoustic parameter set (GeMAPS) for voice research and affective computing. *IEEE Transactions on Affective Computing*, 7(2), 190–202.
- Ezzameli, K., & Mahersia, H. (2023). Emotion recognition from unimodal to multimodal analysis: A review. *Information Fusion*, 99, 101847.
- Fan, Z., Li, M., Zhou, S., & Xu, B. (2020). Exploring wav2vec 2.0 on speaker verification and language identification. *arXiv preprint arXiv:2012.06185*.
- Fayek, H. M., Lech, M., & Cavedon, L. (2015). Towards real-time speech emotion recognition using deep neural networks. In *2015 9th International Conference on Signal Processing and Communication Systems (ICSPCS)* (pp. 1–5).

- Fayek, H. M., Lech, M., & Cavedon, L. (2016). On the Correlation and Transferability of Features Between Automatic Speech Recognition and Speech Emotion Recognition. In *Proceedings of interspeech 2016* (pp. 3618–3622).
- Fayek, H. M., Lech, M., & Cavedon, L. (2017). Evaluating deep learning architectures for speech emotion recognition. *Neural Networks*, *92*, 60–68.
- Feng, H., Ueno, S., & Kawahara, T. (2020). End-to-end speech emotion recognition combined with acoustic-to-word asr model. *Proceedings of Interspeech 2020*, 501–505.
- Feng, T., & Narayanan, S. (2022). Semi-FedSER: Semi-supervised learning for speech emotion recognition on federated learning using multiview pseudo-labeling. In *Proceedings of Interspeech 2022*.
- Feng, T., & Narayanan, S. (2024). Foundation model assisted automatic speech emotion recognition: Transcribing, annotating, and augmenting. In *2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 12116–12120).
- Fernandez, R. (2004). *A computational model for the automatic recognition of affect in speech* (Unpublished doctoral dissertation). Massachusetts Institute of Technology.
- Ferrand, C. T. (2002). Harmonics-to-noise ratio: An index of vocal aging. *Journal of Voice*, *16*(4), 480–487.
- Fleury, A., Noury, N., Vacher, M., Glasson, H., & Seri, J.-F. (2008). Sound and speech detection and classification in a health smart home. In *2008 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (pp. 4644–4647).
- Fontaine, J. R., Scherer, K. R., Roesch, E. B., & Ellsworth, P. C. (2007). The world of emotions is not two-dimensional. *Psychological science*, *18*(12), 1050–1057.
- Fukuda, R., Kano, T., Ando, A., & Ogawa, A. (2025). Speech emotion recognition based on large-scale automatic speech recognizer. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Gao, Y., Chu, C., & Kawahara, T. (2023). Two-stage finetuning of wav2vec 2.0 for speech emotion recognition with asr and gender pretraining. In *Proceedings of Interspeech 2023* (pp. 3637–3641).
- Gao, Y., Wang, L., Liu, J., Dang, J., & Okada, S. (2023). Adversarial domain generalized transformer for cross-corpus speech emotion recognition. *IEEE Transactions on Affective Computing*, *15*(2), 697–708.

- Garofolo, J. S., Lamel, L. F., Fisher, W. M., Fiscus, J. G., & Pallett, D. S. (1993). Darpa timit acoustic-phonetic continuous speech corpus cd-rom. nist speech disc 1-1.1. *NASA STI/Recon technical report n, 93, 27403*.
- Geifman, Y., & El-Yaniv, R. (2017). Selective classification for deep neural networks. In *Advances in neural information processing systems* (Vol. 30).
- Ghosh, S., Rasooli, M. S., Levit, M., Wang, P., Xue, J., Manocha, D., & Li, J. (2024). Failing Forward: Improving Generative Error Correction for ASR with Synthetic Data and Retrieval Augmentation. *arXiv preprint arXiv:2410.13198*.
- Goldwater, S., Jurafsky, D., & Manning, C. D. (2010). Which words are hard to recognize? prosodic, lexical, and disfluency factors that increase speech recognition error rates. *Speech Communication, 52*(3), 181–200.
- Gong, T., Belanich, J., Somandepalli, K., Nagrani, A., Eoff, B., & Jou, B. (2023). LanSER: Language-model supported speech emotion recognition. *Proceedings of Interspeech 2023*.
- Goron, E., Asai, L., Rut, E., & Dinov, M. (2024). Improving domain generalization in speech emotion recognition with whisper. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 11631–11635).
- Grant, K. W., & Greenberg, S. (2001). Speech intelligibility derived from asynchronous processing of auditory-visual information. In *AVSP 2001-International Conference on Auditory-Visual Speech Processing*.
- Graves, A., & Jaitly, N. (2014). Towards end-to-end speech recognition with recurrent neural networks. In *International Conference on Machine Learning* (pp. 1764–1772).
- Gui, A., Gamper, H., Braun, S., & Emmanouilidou, D. (2024). Adapting frechet audio distance for generative music evaluation. In *2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1331–1335).
- Guo, J., Sainath, T. N., & Weiss, R. J. (2019). A spelling correction model for end-to-end speech recognition. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 5651–5655).
- Haider, F., De La Fuente, S., & Luz, S. (2019). An assessment of paralinguistic acoustic features for detection of Alzheimer’s dementia in spontaneous speech. *IEEE Journal of Selected Topics in Signal Processing, 14*(2), 272–281.
- Hama, K., Otsuka, A., & Ishii, R. (2024). Emotion recognition in conversation with multi-step prompting using large language model. In *International Conference on Human-Computer Interaction* (pp. 338–346).

- Hammarberg, B., Fritzell, B., Gauffin, J., Sundberg, J., & Wedin, L. (1980). Perceptual and acoustic correlates of abnormal voice qualities. *Acta Oto-Laryngologica*, 90(1-6), 441–451.
- Han, K., Yu, D., & Tashev, I. (2014). Speech emotion recognition using deep neural network and extreme learning machine. In *Fifteenth Annual Conference of the International Speech Communication Association*.
- Han, W., Jiang, T., Li, Y., Schuller, B., & Ruan, H. (2020). Ordinal learning for emotion recognition in customer service calls. In *ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 6494–6498).
- Han, Z., Geng, T., Feng, H., Yuan, J., Richmond, K., & Li, Y. (2025). Cross-lingual speech emotion recognition: Humans vs. self-supervised models. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Hannun, A., Case, C., Casper, J., Catanzaro, B., Diamos, G., Elsen, E., ... others (2014). Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.
- Hanson, H. M. (1997). Glottal characteristics of female speakers: Acoustic correlates. *The Journal of the Acoustical Society of America*, 101(1), 466–481.
- Hardoon, D. R., Szedmak, S., & Shawe-Taylor, J. (2004). Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12), 2639–2664.
- Hasan, M. K., Lee, S., Rahman, W., Zadeh, A., Mihalcea, R., Morency, L.-P., & Hoque, E. (2021). Humor knowledge enriched transformer for understanding multimodal humor. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 35-14, pp. 12972–12980).
- Hasan, R., Nigar, M., Mamun, N., & Paul, S. (2025). EmoFormer: A Text-Independent Speech Emotion Recognition using a Hybrid Transformer-CNN model. *arXiv preprint arXiv:2501.12682*.
- Hazarika, D., Zimmermann, R., & Poria, S. (2020). MISA: Modality-invariant and-specific representations for multimodal sentiment analysis. In *Proceedings of the 28th ACM International Conference on Multimedia* (pp. 1122–1131).
- He, J., Shi, X., Li, X., & Toda, T. (2024). MF-AED-AEC: Speech Emotion Recognition by Leveraging Multimodal Fusion, ASR Error Detection, and ASR Error Correction. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 11066–11070).

- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
- He, L., Li, Y., Feng, R., Han, X., Liu, Y.-L., Yang, Y., . . . Yuan, J. (2025). Exploring gender bias in alzheimer’s disease detection: Insights from mandarin and greek speech perception. *Proceedings of Interspeech 2025*.
- Hershey, S., Chaudhuri, S., Ellis, D. P., Gemmeke, J. F., Jansen, A., Moore, R. C., . . . others (2017). CNN architectures for large-scale audio classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 131–135).
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8), 2554–2558.
- Hrinchuk, O., Popova, M., & Ginsburg, B. (2020). Correction of automatic speech recognition with transformer sequence-to-sequence model. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 7074–7078).
- Hsu, W.-N., Bolte, B., Tsai, Y.-H. H., Lakhotia, K., Salakhutdinov, R., & Mohamed, A. (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.
- Hu, E. J., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., . . . others (2021). LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Hu, H., Xu, M.-X., & Wu, W. (2007). GMM supervector based SVM with spectral features for speech emotion recognition. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Vol. 4, pp. IV–413).
- Hu, J., Huang, Y., Hu, X., & Xu, Y. (2022). The acoustically emotion-aware conversational agent with speech emotion recognition and empathetic responses. *IEEE Transactions on Affective Computing*, 14(1), 17–30.
- Huang, Z., Dong, M., Mao, Q., & Zhan, Y. (2014). Speech emotion recognition using CNN. In *Proceedings of the 22nd ACM international conference on Multimedia* (pp. 801–804).

- Jacob, A. (2016). Speech emotion recognition based on minimal voice quality features. In *2016 International conference on communication and signal processing (ICCSP)* (pp. 0886–0890).
- Jacob, A. (2017). Modelling speech emotion recognition using logistic regression and decision trees. *International Journal of Speech Technology*, *20*(4), 897–905.
- Jiang, H., Hu, B., Liu, Z., Yan, L., Wang, T., Liu, F., . . . Li, X. (2017). Investigation of different speech types and emotions for detecting depression using different classifiers. *Speech Communication*, *90*, 39–46.
- Jiang, H., Kim, B., Guan, M. Y., & Gupta, M. (2018). To trust or not to trust a classifier. In *Advances in Neural Information Processing Systems* (Vol. 31).
- Jiang, Y., & Poellabauer, C. (2021). A sequence-to-sequence based error correction model for medical automatic speech recognition. In *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* (pp. 3029–3035).
- Kim, B., & Vossen, P. (2021). EmoBERTa: Speaker-aware emotion recognition in conversations with RoBERTa. *arXiv preprint arXiv:2108.12009*.
- Kim, M. J., Yoo, J., Kim, Y., & Kim, H. (2015). Speech emotion classification using tree-structured sparse logistic regression. In *Proceedings of Interspeech 2015* (pp. 1541–1545).
- Koolagudi, S. G., & Rao, K. S. (2012). Emotion recognition from speech using source, system, and prosodic features. *International Journal of Speech Technology*, *15*, 265–289.
- Kumar, P., Kaushik, V., & Raman, B. (2021). Towards the explainability of multimodal speech emotion recognition. In *Proceedings of Interspeech 2021* (pp. 1748–1752).
- Kwon, O.-W., Chan, K., Hao, J., & Lee, T.-W. (2003). Emotion recognition by speech signals. In *Proceedings of Interspeech 2003* (pp. 125–128).
- Lai, C.-I., Chuang, Y.-S., Lee, H.-Y., Li, S.-W., & Glass, J. (2021). Semi-supervised spoken language understanding via self-supervised speech and language model pretraining. In *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 7468–7472).
- Lakomkin, E., Weber, C., Magg, S., & Wermter, S. (2017). Reusing neural speech representations for auditory emotion recognition. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (pp. 423–430).
- Lakomkin, E., Weber, C., Magg, S., & Wermter, S. (2018). Reusing neural speech representations for auditory emotion recognition. *arXiv preprint arXiv:1803.11508*.

- Lakomkin, E., Zamani, M. A., Weber, C., Magg, S., & Wermter, S. (2018). On the robustness of speech emotion recognition for human-robot interaction with deep neural networks. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 854–860).
- Lanjewar, R. B., Mathurkar, S., & Patel, N. (2015). Implementation and comparison of speech emotion recognition system using Gaussian Mixture Model (GMM) and K-Nearest Neighbor (K-NN) techniques. *Procedia computer science*, *49*, 50–57.
- Laplante, D., & Ambady, N. (2003). On how things are said: Voice tone, voice intensity, verbal content, and perceptions of politeness. *Journal of language and social psychology*, *22*(4), 434–441.
- Latif, S., Usama, M., Malik, M. I., & Schuller, B. W. (2023). Can large language models aid in annotating speech emotional data? uncovering new frontiers. *arXiv preprint arXiv:2307.06090*.
- LeCun, Y., & Bengio, Y. (1995). Convolutional networks for images, speech, and time series. *The Handbook of Brain Theory and Neural Networks*, *3361*, 255–258.
- Lee, C. M., & Narayanan, S. S. (2005). Toward detecting emotions in spoken dialogs. *IEEE Transactions on Speech and Audio Processing*, *13*(2), 293–303.
- Lee, J., & Tashev, I. (2015). High-level feature representation using recurrent neural network for speech emotion recognition. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- Leem, S.-G., Fulford, D., Onnela, J.-P., Gard, D., & Busso, C. (2023). Selective acoustic feature enhancement for speech emotion recognition with noisy speech. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, *32*, 917–929.
- Levin, K., Henry, K., Jansen, A., & Livescu, K. (2013). Fixed-dimensional acoustic embeddings of variable-length segments in low-resource settings. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding* (pp. 410–415).
- Li, B., Dimitriadis, D., & Stolcke, A. (2019). Acoustic and lexical sentiment analysis for customer service calls. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 5876–5880).
- Li, C., Wang, J., Zhang, Y., Zhu, K., Hou, W., Lian, J., ... Xie, X. (2023). Large language models understand and can be enhanced by emotional stimuli. *arXiv preprint arXiv:2307.11760*.
- Li, C., Xu, W., Cohen, T., & Pakhomov, S. (2024). Useful blunders: Can automated speech recognition errors improve downstream dementia classification? *Journal of Biomedical Informatics*, *150*, 104598.

- Li, H., Ding, W., Wu, Z., & Liu, Z. (2020). Learning fine-grained cross modality excitement for speech emotion recognition. *Proceedings of Interspeech 2020*.
- Li, J., Lavrukhin, V., Ginsburg, B., Leary, R., Kuchaiev, O., Cohen, J. M., ... Gadde, R. T. (2019). Jasper: An end-to-end convolutional neural acoustic model. *arXiv preprint arXiv:1904.03288*.
- Li, M., Yang, D., Liu, Y., Wang, S., Chen, J., Wang, S., ... others (2024). Toward robust incomplete multimodal sentiment analysis via hierarchical representation learning. *arXiv preprint arXiv:2411.02793*.
- Li, M., Yang, D., Zhao, X., Wang, S., Wang, Y., Yang, K., ... Zhang, L. (2024). Correlation-Decoupled Knowledge Distillation for Multimodal Sentiment Analysis with Incomplete Modalities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 12458–12468).
- Li, Q., Qiu, D., Zhang, Y., Li, B., He, Y., Woodland, P. C., ... Strohmman, T. (2021). Confidence estimation for attention-based sequence-to-sequence models for speech recognition. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6388–6392).
- Li, Q., Zhang, Y., Qiu, D., He, Y., Cao, L., & Woodland, P. C. (2022). Improving confidence estimation on out-of-domain data for end-to-end speech recognition. In *Icassp 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (icassp)* (pp. 6537–6541).
- Li, X., Tao, J., Johnson, M. T., Soltis, J., Savage, A., Leong, K. M., & Newman, J. D. (2007). Stress and emotion classification using jitter and shimmer features. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Vol. 4, pp. IV–1081).
- Li, Y. (2018). Towards improving speech emotion recognition for in-vehicle agents: Preliminary results of incorporating sentiment analysis by using early and late fusion methods. In *Proceedings of the 6th International Conference on Human-Agent Interaction* (pp. 365–367).
- Li, Y. (2021a, August 31). *Feeling estimation device, feeling estimation method, and storage medium*. Google Patents. (US Patent 11,107,464)
- Li, Y. (2021b). Semi-supervised learning for multimodal speech and emotion recognition. In *Proceedings of the 2021 International Conference on Multimodal Interaction* (pp. 817–821).
- Li, Y. (2022, September 13). *Information processing apparatus, information processing method, and storage medium*. Google Patents. (US Patent 11,443,759)

- Li, Y. (2023a). Enhancing Speech Emotion Recognition for Real-World Applications via ASR Integration. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)* (pp. 1–5).
- Li, Y. (2023b, July 25). *Information-processing device, vehicle, computer-readable storage medium, and information-processing method*. Google Patents. (US Patent 11,710,499)
- Li, Y., Bell, P., & Lai, C. (2022). Fusing ASR outputs in joint training for speech emotion recognition. In *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 7362–7366).
- Li, Y., Bell, P., & Lai, C. (2023a). Multimodal dyadic impression recognition via listener adaptive cross-domain fusion. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Li, Y., Bell, P., & Lai, C. (2023b). Transfer learning for personality perception via speech emotion recognition. In *Proceedings of interspeech 2023* (pp. 5197–5201).
- Li, Y., Bell, P., & Lai, C. (2024). Speech emotion recognition with ASR transcripts: A comprehensive study on word error rate and fusion techniques. *2024 IEEE Spoken Language Technology Workshop (SLT)*. IEEE.
- Li, Y., Chen, P., Bell, P., & Lai, C. (2024). Crossmodal ASR error correction with discrete speech units. *2024 IEEE Spoken Language Technology Workshop (SLT)*. IEEE.
- Li, Y., Gong, Y., Yang, C.-H. H., Bell, P., & Lai, C. (2025). Revise, Reason, and Recognize: LLM-Based Emotion Recognition via Emotion-Specific Prompts and ASR Error Correction. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Li, Y., Gui, A., Emmanouilidou, D., & Gamper, H. (2025). Addressing emotion bias in music emotion recognition and generation with frechet audio distance. *2025 IEEE International Conference on Multimedia and Expo (ICME)*.
- Li, Y., Inoue, K., Tian, L., Fu, C., Ishi, C. T., Ishiguro, H., . . . Lai, C. (2023). I know your feelings before you do: Predicting future affective reactions in human-computer dialogue. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–7).
- Li, Y., Ishi, C. T., Inoue, K., Nakamura, S., & Kawahara, T. (2019). Expressing reactive emotion based on multimodal emotion recognition for natural conversation in human–robot interaction. *Advanced Robotics*, *33*(20), 1030–1041.

- Li, Y., Kollias, D., Chanel, G., Fanourakis, M., Muszynski, M., M. Booth, B., . . . Chen, H. (2025). Hrai 2025: The 1st workshop on holistic and responsible affective intelligence. In *Proceedings of the 27th international conference on multimodal interaction* (pp. 814–817).
- Li, Y., & Lai, C. (2022a). A cross-domain approach for continuous impression recognition from dyadic audio-visual-physio signals. *arXiv preprint arXiv:2203.13932*.
- Li, Y., & Lai, C. (2022b). Robotic speech synthesis: Perspectives on interactions, scenarios, and ethics. *arXiv preprint arXiv:2203.09599*.
- Li, Y., & Lai, C. (2023). Empowering dialogue systems with affective and adaptive interaction: Integrating social intelligence. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)* (pp. 1–8).
- Li, Y., Lai, C., Lala, D., Inoue, K., & Kawahara, T. (2022). Alzheimer’s Dementia Detection through Spontaneous Dialogue with Proactive Robotic Listeners. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 875–879).
- Li, Y., Mohamied, Y., Bell, P., & Lai, C. (2023). Exploration of a self-supervised speech model: A study on emotional corpora. In *2022 IEEE Spoken Language Technology Workshop (SLT)* (pp. 868–875).
- Li, Y., Urquhart, L., Karatas, N., Shao, S., Ishiguro, H., & Shen, X. (2024). Beyond voice assistants: Exploring advantages and risks of an in-car social robot in real driving scenarios. *arXiv preprint arXiv:2402.11853*.
- Li, Y., Wang, Y., & Cui, Z. (2023). Decoupled multimodal distilling for emotion recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 6631–6640).
- Li, Y., Williams, J., Feng, T., Mitra, V., Gong, Y., Shi, B., . . . Bell, P. (2023). Responsible speech foundation models. *Interspeech*.
- Li, Y., Zhang, Z., Han, J., Bell, P., & Lai, C. (2025). Semi-Supervised Cognitive State Classification from Speech with Multi-View Pseudo-Labeling. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Li, Y., Zhao, T., Kawahara, T., et al. (2019). Improved end-to-end speech emotion recognition using self attention mechanism and multitask learning. In *Interspeech* (pp. 2803–2807).
- Li, Y., Zhao, T., & Shen, X. (2020). Attention-Based Multimodal Fusion for Estimating Human Emotion in Real-World HRI. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 340–342).

- Li, Y., Zhao, Z., Klejch, O., Bell, P., & Lai, C. (2023). ASR and emotional speech: A word-level investigation of the mutual impact of speech and emotion recognition. In *Proceedings of interspeech 2023*.
- Lian, Z., Chen, L., Sun, L., Liu, B., & Tao, J. (2023). GCNet: Graph completion network for incomplete multimodal learning in conversation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7), 8419–8432.
- Lian, Z., Sun, H., Sun, L., Wen, Z., Zhang, S., Chen, S., . . . others (2024). MER 2024: Semi-supervised learning, noise robustness, and open-vocabulary multimodal emotion recognition. In *Proceedings of the 2nd International Workshop on Multimodal and Responsible Affective Computing* (pp. 41–48).
- Liang, P. P., Liu, Z., Zadeh, A., & Morency, L.-P. (2018). Multimodal Language Analysis with Recurrent Multistage Fusion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (pp. 150–161).
- Liao, J., Eskimez, S., Lu, L., Shi, Y., Gong, M., Shou, L., . . . Zeng, M. (2023). Improving readability for automatic speech recognition transcription. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 22(5), 1–23.
- Lin, B., & Wang, L. (2023). Multi-modal ASR error correction with joint ASR error detection. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Lin, Y.-L., & Wei, G. (2005). Speech emotion recognition based on HMM and SVM. In *2005 International Conference on Machine Learning and Cybernetics* (Vol. 8, pp. 4898–4901).
- Littlewort, G., Whitehill, J., Wu, T., Fasel, I., Frank, M., Movellan, J., & Bartlett, M. (2011). The computer expression recognition toolbox (CERT). In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)* (pp. 298–305).
- Liu, J., Chen, C., Bu, J., You, M., & Tao, J. (2007). Speech emotion recognition using an enhanced co-training algorithm. In *2007 IEEE International Conference on Multimedia and Expo* (pp. 999–1002).
- Liu, Y. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Liu, Y.-L., Li, Y., Feng, R., He, L., Chen, J.-X., Wang, Y.-M., . . . Ling, Z.-H. (2025). Leveraging cascaded binary classification and multimodal fusion for dementia detection through spontaneous speech. *arXiv preprint arXiv:2505.19446*.

- Liu, Z., Shen, Y., Lakshminarasimhan, V. B., Liang, P. P., Zadeh, A. B., & Morency, L.-P. (2018). Efficient Low-rank Multimodal Fusion With Modality-Specific Factors. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 2247–2256).
- Liu, Z., Zhou, B., Chu, D., Sun, Y., & Meng, L. (2024). Modality translation-based multimodal sentiment analysis under uncertain missing modalities. *Information Fusion, 101*, 101973.
- Livingstone, S. R., & Russo, F. A. (2018). The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS): A dynamic, multimodal set of facial and vocal expressions in North American English. *PloS one, 13*(5), e0196391.
- Lotfian, R., & Busso, C. (2017). Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings. *IEEE Transactions on Affective Computing, 10*(4), 471–483.
- Lu, J., Batra, D., Parikh, D., & Lee, S. (2019). ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in neural information processing systems*.
- Lu, K., Foo, C.-S., Teh, K. K., Tran, H. D., & Chandrasekhar, V. R. (2019). Semi-Supervised Audio Classification with Consistency-Based Regularization. In *Proceedings of Interspeech 2019* (Vol. 1, pp. 3654–3658).
- Luengo, I., Navas, E., & Hernáez, I. (2010). Feature analysis and evaluation for automatic emotion identification in speech. *IEEE Transactions on Multimedia, 12*(6), 490–501.
- Luengo, I., Navas, E., Hernáez, I., & Sánchez, J. (2005). Automatic emotion recognition using prosodic parameters. In *Proceedings of Interspeech 2005* (pp. 493–496).
- Lugger, M., & Yang, B. (2007). The relevance of voice quality features in speaker independent emotion recognition. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Vol. 4, pp. IV–17).
- Luo, J., Phan, H., & Reiss, J. (2023). Fine-tuned RoBERTa Model with a CNN-LSTM Network for Conversational Emotion Recognition. In *Proceedings of Interspeech 2023* (pp. 2413–2417).
- Luo, W., Xu, M., & Lai, H. (2023). Multimodal reconstruct and align net for missing modality problem in sentiment analysis. In *International Conference on Multimedia Modeling* (pp. 411–422).
- Luz, S., Haider, F., de la Fuente, S., Fromm, D., & MacWhinney, B. (2021). Detecting cognitive decline using speech only: The ADReSSo Challenge. In *Proceedings of Interspeech*.

- Ma, F., Zhang, Y., & Sun, X. (2023). Multimodal Sentiment Analysis with Preferential Fusion and Distance-aware Contrastive Learning. In *2023 IEEE International Conference on Multimedia and Expo (ICME)* (pp. 1367–1372).
- Ma, Z., Zheng, Z., Ye, J., Li, J., Gao, Z., Zhang, S., & Chen, X. (2024). emotion2vec: Self-Supervised Pre-Training for Speech Emotion Representation. In *Findings of the Association for Computational Linguistics ACL 2024* (pp. 15747–15760).
- Maas, A. L., Miller, S. D., O'neil, T. M., Ng, A. Y., & Nguyen, P. (2012). Word-level acoustic modeling with convolutional vector regression. In *Proceedings of ICML Workshop Representation Learning*.
- Mani, A., Palaskar, S., Meripo, N. V., Konam, S., & Metze, F. (2020). ASR error correction and domain adaptation using machine translation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6344–6348).
- Marcus, M., Santorini, B., & Marcinkiewicz, M. A. (1993). Building a large annotated corpus of english: The penn treebank. *Computational Linguistics*, 19(2), 313–330.
- Matusevych, Y., Kamper, H., & Goldwater, S. (2020). Analyzing autoencoder-based acoustic word embeddings. In *ICLR Workshop on Bridging AI and Cognitive Science*.
- McAuliffe, M., Socolof, M., Mihuc, S., Wagner, M., & Sonderegger, M. (2017). Montreal forced aligner: Trainable text-speech alignment using Kaldi. In *Proceedings of Interspeech 2017* (Vol. 2017, pp. 498–502).
- Mirsamadi, S., Barsoum, E., & Zhang, C. (2017). Automatic speech emotion recognition using recurrent neural networks with local attention. In *2017 IEEE International conference on acoustics, speech and signal processing (ICASSP)* (pp. 2227–2231).
- Moell, B., Aronsson, F. S., Östberg, P., & Beskow, J. (2025). The order in speech disorder: a scoping review of state of the art machine learning methods for clinical speech classification. *arXiv preprint arXiv:2503.04802*.
- Moore, B. C. (2012). *An introduction to the psychology of hearing*. Brill.
- Morgan, S. D. (2019). Categorical and dimensional ratings of emotional speech: Behavioral findings from the morgan emotional speech set. *Journal of Speech, Language, and Hearing Research*, 62(11), 4015–4029.
- Mote, P., Sisman, B., & Busso, C. (2024). Unsupervised domain adaptation for speech emotion recognition using k-nearest neighbors voice conversion. In *Proceedings of Interspeech 2024*.

- Munot, R., & Nenkova, A. (2019). Emotion impacts speech recognition performance. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop* (pp. 16–21).
- Noroozi, F., Sapiński, T., Kamińska, D., & Anbarjafari, G. (2017). Vocal-based emotion recognition using random forests and decision tree. *International Journal of Speech Technology*, 20(2), 239–246.
- Nwe, T. L., Foo, S. W., & De Silva, L. C. (2003). Speech emotion recognition using hidden markov models. *Speech Communication*, 41(4), 603–623.
- Padi, T., Singh, A., & Lal, S. (2022). Bert-based multimodal emotion recognition using acoustic and text features. *arXiv preprint arXiv:2202.08974*.
- Pan, Y., Mirheidari, B., Harris, J. M., Thompson, J. C., Jones, M., Snowden, J. S., ... Christensen, H. (2021). Using the Outputs of Different Automatic Speech Recognition Paradigms for Acoustic-and BERT-Based Alzheimer's Dementia Detection Through Spontaneous Speech. In *Proceedings of Interspeech 2021* (pp. 3810–3814).
- Pan, Y., Mirheidari, B., Reuber, M., Venneri, A., Blackburn, D., & Christensen, H. (2020). Improving detection of alzheimer's disease using automatic speech recognition to identify high-quality segments for more robust feature extraction. In *Proceedings of interspeech 2020* (pp. 4961–4965).
- Panayotov, V., Chen, G., Povey, D., & Khudanpur, S. (2015). Librispeech: an ASR corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 5206–5210).
- Papineni, K., Roukos, S., Ward, T., & Zhu, W.-J. (2002). BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics* (pp. 311–318).
- Park, S., Kim, J., Ye, S., Jeon, J., Park, H. Y., & Oh, A. (2021). Dimensional Emotion Detection from Categorical Emotion. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing* (pp. 4367–4380).
- Pasad, A., Chou, J.-C., & Livescu, K. (2021). Layer-wise Analysis of a Self-supervised Speech Representation Model. In *Proceedings of ASRU 2021*.
- Pasad, A., Shi, B., & Livescu, K. (2023). Comparative layer-wise analysis of self-supervised speech models. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Pascual, S., Ravanelli, M., Serra, J., Bonafonte, A., & Bengio, Y. (2019). Learning problem-agnostic speech representations from multiple self-supervised tasks. In *Proceedings of Interspeech 2019*.

- Pastor, M. A., Ribas, D., Ortega, A., Miguel, A., & Lleida, E. (2022). Cross-corpus speech emotion recognition with hubert self-supervised representation. In *IberSPEECH 2022* (pp. 76–80).
- Peng, Z., Lu, Y., Pan, S., & Liu, Y. (2021). Efficient speech emotion recognition using multi-scale CNN and attention. In *ICASSP 2021-2021 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 3020–3024).
- Peng, Z., Zhu, Z., Unoki, M., Dang, J., & Akagi, M. (2018). Auditory-inspired end-to-end speech emotion recognition using 3D convolutional recurrent neural networks based on spectral-temporal representation. In *2018 IEEE International Conference on Multimedia and Expo (ICME)* (pp. 1–6).
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (emnlp)* (pp. 1532–1543).
- Pepino, L., Riera, P., & Ferrer, L. (2021). Emotion Recognition from Speech Using wav2vec 2.0 Embeddings. In *Proceedings of Interspeech 2021* (pp. 3400–3404).
- Pepino, L., Riera, P., Ferrer, L., & Gravano, A. (2020). Fusion approaches for emotion recognition from speech using acoustic and text-based features. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6484–6488).
- Peterson, G. E., & Barney, H. L. (1952). Control methods used in a study of the vowels. *The Journal of the Acoustical Society of America*, 24(2), 175–184.
- Pham, H., Liang, P. P., Manzini, T., Morency, L.-P., & Póczos, B. (2019). Found in translation: Learning robust joint representations by cyclic translations between modalities. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 33, pp. 6892–6899).
- Pohjalainen, J., Fabien Ringeval, F., Zhang, Z., & Schuller, B. (2016). Spectral and cepstral audio noise reduction techniques in speech emotion recognition. In *Proceedings of the 24th ACM international Conference on Multimedia* (pp. 670–674).
- Polzehl, T., Schmitt, A., Metze, F., & Wagner, M. (2011). Anger recognition in speech using acoustic and linguistic cues. *Speech Communication*, 53(9-10), 1198–1209.
- Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., ... others (2011). The kaldi speech recognition toolkit. In *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*.
- Qiu, D., Li, Q., He, Y., Zhang, Y., Li, B., Cao, L., ... others (2021). Learning word-level confidence for subword end-to-end ASR. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6393–6397).

- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., . . . others (2021). Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning* (pp. 8748–8763).
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., & Sutskever, I. (2023). Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning* (pp. 28492–28518).
- Radhakrishnan, S., Yang, C.-H. H., Khan, S. A., Kumar, R., Kiani, N. A., Gomez-Cabrero, D., & Tegnér, J. (2023). Whispering LLaMA: A Cross-Modal Generative Error Correction Framework for Speech Recognition. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Rahman, W., Hasan, M. K., Lee, S., Zadeh, A., Mao, C., Morency, L.-P., & Hoque, E. (2020). Integrating multimodal information in large pretrained transformers. In *Proceedings of the conference. Association for Computational Linguistics. Meeting* (Vol. 2020, p. 2359).
- Rao, K. S., Koolagudi, S. G., & Vempada, R. R. (2013). Emotion recognition from speech using global and local prosodic features. *International Journal of Speech Technology*, 16, 143–160.
- Rashed, A., Elsayed, S., & Schmidt-Thieme, L. (2022). Context and Attribute-Aware Sequential Recommendation via Cross-Attention. In *Proceedings of the 16th ACM Conference on Recommender Systems* (pp. 71–80).
- Rezaie, V., Parnianifard, A., Zegarra Rodriguez, D., Mumtaz, S., & Wuttisittikulij, L. (2022). Speech Emotion Recognition using ANFIS and PSO-optimization with Word2Vec. *Research Square*.
- Roblek, D., Kilgour, K., Sharifi, M., & Zuluaga, M. (2019). Fréchet Audio Distance: A Reference-free Metric for Evaluating Music Enhancement Algorithms. In *Proceedings of Interspeech 2019* (pp. 2350–2354).
- Rozgić, V., Ananthakrishnan, S., Saleem, S., & Hansen, J. H. (2012). Emotion recognition using acoustic and lexical features. In *Proceedings of Interspeech 2012* (pp. 199–202).
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533–536.
- Russell, J. A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6), 1161.
- Sahay, S., Okur, E., Kumar, S. H., & Nachman, L. (2020). Low rank fusion based transformers for multimodal sequences. In *Second Grand-Challenge and Workshop on Multimodal Language (Challenge-HML)* (pp. 29–34).

- Sahu, S., Mitra, V., Seneviratne, N., & Espy-Wilson, C. (2019b). Multi-Modal Learning for Speech Emotion Recognition: An Analysis and Comparison of ASR Outputs with Ground Truth Transcription. In *Proceedings of Interspeech 2019* (pp. 3302–3306).
- Sahu, S., Mitra, V., Seneviratne, N., & Espy-Wilson, C. Y. (2019a). Multi-Modal Learning for Speech Emotion Recognition: An Analysis and Comparison of ASR Outputs with Ground Truth Transcription. In *Proceedings of interspeech 2019* (pp. 3302–3306).
- Saliba, A., Li, Y., Sanabria, R., & Lai, C. (2024). Layer-Wise Analysis of Self-Supervised Acoustic Word Embeddings: A Study on Speech Emotion Recognition. In *2024 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*.
- Sanabria, R., Tang, H., & Goldwater, S. (2023). Analyzing acoustic word embeddings from pre-trained self-supervised speech models. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Sanders, N., Li, Y., Richmond, K., & King, S. (2025). Segmentation-variant codebooks for preservation of paralinguistic and prosodic information. *Proceedings of Interspeech 2025*.
- Santoso, J., Ishizuka, K., & Hashimoto, T. (2024). Large language model-based emotional speech annotation using context and acoustic feature for speech emotion recognition. In *2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 11026–11030).
- Santoso, J., Yamada, T., Ishizuka, K., Hashimoto, T., & Makino, S. (2022). Speech emotion recognition based on self-attention weight correction for acoustic and text features. *IEEE Access*, *10*, 115732–115743.
- Santoso, J., Yamada, T., Makino, S., Ishizuka, K., & Hiramura, T. (2021). Speech Emotion Recognition Based on Attention Weight Correction Using Word-Level Confidence Measure. In *Proceedings of Interspeech 2021* (pp. 1947–1951).
- Sap, M., Rashkin, H., Chen, D., Le Bras, R., & Choi, Y. (2019). Social IQa: Commonsense reasoning about social interactions. In *EMNLP-IJCNLP* (pp. 4463–4473).
- Satt, A., Rozenberg, S., & Hoory, R. (2017). Efficient emotion recognition from speech using deep learning on spectrograms. In *Proceedings of Interspeech 2017* (pp. 1089–1093).
- Savchenko, A. V. (2022). HSEmotion: High-speed emotion recognition library. *Software Impacts*, *14*, 100433.
- Schuller, B., & Batliner, A. (2013). *Computational paralinguistics: emotion, affect and personality in speech and language processing*. John Wiley & Sons.

- Schuller, B., Batliner, A., Steidl, S., & Seppi, D. (2009). Emotion recognition from speech: putting asr in the loop. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 4585–4588).
- Schuller, B., Steidl, S., & Batliner, A. (2009). The Interspeech 2009 emotion challenge. In *Tenth Annual Conference of the International Speech Communication Association*.
- Schuller, B., Steidl, S., Batliner, A., Vinciarelli, A., Scherer, K., Ringeval, F., ... others (2013). The INTERSPEECH 2013 computational paralinguistics challenge: social signals, conflict, emotion, autism. In *Proceedings interspeech 2013*.
- Schuller, B., Vlasenko, B., Eyben, F., Rigoll, G., & Wendemuth, A. (2009). Acoustic emotion recognition: A benchmark comparison of performances. In *2009 IEEE Workshop on Automatic Speech Recognition & Understanding* (pp. 552–557).
- Sciar, M., Choi, Y., Tsvetkov, Y., & Suhr, A. (2024). Quantifying Language Models' Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting. In *International Conference on Learning Representations (ICLR)*.
- Scotti, V., Galati, F., Sbattella, L., & Tedesco, R. (2021). Combining deep and unsupervised features for multilingual speech emotion recognition. In *International Conference on Pattern Recognition* (pp. 114–128).
- Sebastian, J., Pierucci, P., et al. (2019). Fusion Techniques for Utterance-Level Emotion Recognition Combining Speech and Transcripts. In *Proceedings of Interspeech 2019* (pp. 51–55).
- Sharma, M. (2022). Multi-Lingual Multi-Task Speech Emotion Recognition Using wav2vec 2.0. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6907–6911).
- Shen, P., Changjun, Z., & Chen, X. (2011). Automatic speech emotion recognition using support vector machine. In *Proceedings of 2011 International Conference on Electronic & Mechanical Engineering and Information Technology* (Vol. 2, pp. 621–625).
- Shen, X., Luo, Z., Li, Y., Ouyang, T., & Wu, Y. (2024). Chance-constrained abnormal data cleaning for robust classification with noisy labels. *IEEE Transactions on Emerging Topics in Computational Intelligence*.
- Shon, S., Brusco, P., Pan, J., Han, K. J., & Watanabe, S. (2021). Leveraging Pre-trained Language Model for Speech Sentiment Analysis. In *Proceedings of Interspeech 2021* (pp. 566–570).
- Singh, S., Caragea, C., & Li, J. J. (2024). Language models (mostly) do not consider emotion triggers when predicting emotion. *NAACL 2024*.

- Siriwardhana, S., Reis, A., Weerasekera, R., & Nanayakkara, S. (2020). Jointly Fine-Tuning “BERT-Like” Self Supervised Models to Improve Multimodal Speech Emotion Recognition. In *Proceedings of Interspeech 2020* (pp. 3755–3759).
- Snoek, C. G., Worring, M., & Smeulders, A. W. (2005). Early versus late fusion in semantic video analysis. In *Proceedings of the 13th Annual ACM International Conference on Multimedia* (pp. 399–402).
- Sobin, C., & Alpert, M. (1999). Emotion in speech: The acoustic attributes of fear, anger, sadness, and joy. *Journal of Psycholinguistic Research*, 28, 347–365.
- Soltau, H., Liao, H., & Sak, H. (2017). Neural Speech Recognizer: Acoustic-to-Word LSTM Model for Large Vocabulary Speech Recognition. In *Proceedings of Interspeech 2017* (pp. 3707–3711).
- Stevens, K. N. (1998). *Acoustic phonetics*. MIT Press.
- Sun, Y., Zhao, Z., Richmond, K., & Li, Y. (2025). Exploring Acoustic Similarity in Emotional Speech and Music via Self-Supervised Representations. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Tanaka, T., Masumura, R., Ihori, M., Takashima, A., Moriya, T., Ashihara, T., . . . Makishima, N. (2021). Cross-modal transformer-based neural correction models for automatic speech recognition. *Proceedings of Interspeech 2021*.
- Tanaka, T., Masumura, R., Masataki, H., & Aono, Y. (2018). Neural Error Corrective Language Models for Automatic Speech Recognition. In *Proceedings of Interspeech 2018* (pp. 401–405).
- Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., . . . Hashimoto, T. B. (2023). *Stanford alpaca: An instruction-following llama model*.
- Tian, L., Moore, J., & Lai, C. (2016). Recognizing emotions in spoken dialogue with hierarchically fused acoustic and lexical features. In *Proceedings of SLT 2016* (pp. 565–572).
- Tits, N., El Haddad, K., & Dutoit, T. (2018). ASR-based Features for Emotion Recognition: A Transfer Learning Approach. In *Proceedings of Grand Challenge and Workshop on Human Multimodal Language (Challenge-HML)* (pp. 48–52).
- Titze, I. R. (1994). *Principles of voice production*. Prentice Hall.
- Tokuhisa, R., Inui, K., & Matsumoto, Y. (2008). Emotion classification using massive examples extracted from the web. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)* (pp. 881–888).

- Tompkins, D., Emmanouilidou, D., Deshmukh, S., & Elizalde, B. (2023). Multi-view learning for speech emotion recognition with categorical emotion, categorical sentiment, and dimensional scores. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5).
- Triantafyllopoulos, A., Wagner, J., Wierstorf, H., Schmitt, M., Reichel, U., Eyben, F., ... Schuller, B. W. (2022). Probing speech emotion recognition transformers for linguistic knowledge. In *Proceedings of Interspeech 2022* (pp. 146–150).
- Tsai, Y.-H. H., Bai, S., Liang, P. P., Kolter, J. Z., Morency, L.-P., & Salakhutdinov, R. (2019). Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting* (Vol. 2019, p. 6558).
- Tsai, Y.-H. H., Liang, P. P., Zadeh, A., Morency, L.-P., & Salakhutdinov, R. (2019). Learning Factorized Multimodal Representations. In *International Conference on Representation Learning*.
- Tzirakis, P., Zhang, J., & Schuller, B. W. (2018). End-to-end speech emotion recognition using deep neural networks. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 5089–5093).
- Upadhyay, S. G., Busso, C., & Lee, C.-C. (2024). A Layer-Anchoring Strategy for Enhancing Cross-Lingual Speech Emotion Recognition. In *Proceedings of Interspeech 2024* (pp. 4693–4697).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems* (pp. 5998–6008).
- Venkata Subbarao, M., Terlapu, S. K., Geethika, N., & Harika, K. D. (2021). Speech emotion recognition using k-nearest neighbor classifiers. In *Recent Advances in Artificial Intelligence and Data Engineering: Select Proceedings of AIDE 2020* (pp. 123–131). Springer.
- Wagner, J., Triantafyllopoulos, A., Wierstorf, H., Schmitt, M., Burkhardt, F., Eyben, F., & Schuller, B. W. (2023). Dawn of the transformer era in speech emotion recognition: closing the valence gap. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wallington, E., Kershenbaum, B., Bell, P., & Klejch, O. (2021). On the learning dynamics of semi-supervised training for ASR. In *Proceedings of Interspeech 2021* (pp. 716–720).
- Wang, K., An, N., Li, B. N., Zhang, Y., & Li, L. (2015). Speech emotion recognition using fourier parameters. *IEEE Transactions on Affective Computing*, 6(1), 69–75.

- Wang, X., Girshick, R., Gupta, A., & He, K. (2018). Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 7794–7803).
- Wang, X., Li, C., Chang, Y., Wang, J., & Wu, Y. (2024). Negativeprompt: Leveraging psychology for large language models enhancement via negative emotional stimuli. *IJCAI 2024*.
- Wang, X., Li, X., Yin, Z., Wu, Y., & Liu, J. (2023). Emotional intelligence of large language models. *Journal of Pacific Rim Psychology*.
- Wang, X., Wang, M., Qi, W., Su, W., Wang, X., & Zhou, H. (2021). A novel end-to-end speech emotion recognition network with stacked transformer layers. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6289–6293).
- Wang, Y., Li, Y., Liang, P. P., Morency, L.-P., Bell, P., & Lai, C. (2023). Cross-attention is not enough: Incongruity-aware dynamic hierarchical fusion for multimodal affect recognition. *arXiv preprint arXiv:2305.13583*.
- Wang, Y., Shen, G., Xu, Y., Li, J., & Zhao, Z. (2021). Learning Mutual Correlation in Multimodal Transformer for Speech Emotion Recognition. In *Proceedings of Interspeech 2021* (pp. 4518–4522).
- Wang, Y., Shen, Y., Liu, Z., Liang, P. P., Zadeh, A., & Morency, L.-P. (2019). Words can shift: Dynamically adjusting word representations using nonverbal behaviors. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 33, pp. 7216–7223).
- Wang, Z., Wan, Z., & Wan, X. (2020). Transmodality: An end2end fusion method with transformer for multimodal sentiment analysis. In *Proceedings of the Web Conference 2020* (pp. 2514–2520).
- Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior Research Methods*, 45, 1191–1207.
- Watanabe, S., Hori, T., Karita, S., Hayashi, T., Nishitoba, J., Unno, Y., ... others (2018). Espnet: End-to-end speech processing toolkit. *Proceedings of Interspeech 2018*.
- Wu, S., Falk, T. H., & Chan, W.-Y. (2011). Automatic speech emotion recognition using modulation spectral features. *Speech Communication*, 53(5), 768–785.
- Wu, W., Zhang, C., & Woodland, P. C. (2021). Emotion recognition by fusing time synchronous and time asynchronous representations. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6269–6273).

- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., ... others (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Wu, Y., Zhao, Y., Lu, X., Qin, B., Wu, Y., Sheng, J., & Li, J. (2021). Modeling incongruity between modalities for multimodal sarcasm detection. *IEEE MultiMedia*, 28(2), 86–95.
- Xie, Q., Luong, M.-T., Hovy, E., & Le, Q. V. (2020). Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 10687–10698).
- Xu, H., Zhang, H., Han, K., Wang, Y., Peng, Y., & Li, X. (2019). Learning alignment for multimodal emotion recognition from speech. In *Proceedings of Interspeech 2019* (pp. 3569–3573).
- Xu, M., Zhang, F., & Zhang, W. (2021). Head fusion: Improving the accuracy and robustness of speech emotion recognition on the iemocap and ravedss dataset. *IEEE Access*, 9, 74539–74549.
- Xu, N., Mao, W., & Chen, G. (2018). A co-memory network for multimodal sentiment analysis. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (pp. 929–932).
- Yang, C.-H. H., Gu, Y., Liu, Y.-C., Ghosh, S., Bulyko, I., & Stolcke, A. (2023). Generative speech recognition error correction with large language models and task-activating prompting. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)* (pp. 1–8).
- Yang, C.-H. H., Park, T., Gong, Y., Li, Y., Chen, Z., Lin, Y.-T., ... others (2024). Large language model based generative error correction: A challenge and baselines for speech recognition, speaker tagging, and emotion recognition. In *2024 IEEE Spoken Language Technology Workshop (SLT)* (pp. 371–378).
- Yang, J., Liu, J., Huang, K., Xia, J., Zhu, Z., & Zhang, H. (2024). Single-and cross-lingual speech emotion recognition based on wavlm domain emotion embedding. *Electronics*, 13(7), 1380.
- Yang, K., Xu, H., & Gao, K. (2020). Cm-bert: Cross-modal bert for text-audio sentiment analysis. In *Proceedings of the 28th ACM International Conference on Multimedia* (pp. 521–528).
- Yang, S.-w., Chi, P.-H., Chuang, Y.-S., Lai, C.-I. J., Lakhotia, K., Lin, Y. Y., ... others (2021). SUPERB: Speech Processing Universal PERFORMANCE Benchmark. In *Proceedings of Interspeech 2021* (pp. 1194–1198).

- Yeh, S.-L., Lin, Y.-S., & Lee, C.-C. (2020). Speech Representation Learning for Emotion Recognition Using End-to-End ASR with Factorized Adaptation. In *Proceedings of Interspeech 2020* (pp. 536–540).
- Yenigalla, P., Kumar, A., Tripathi, S., Singh, C., Kar, S., & Vepa, J. (2018). Speech Emotion Recognition Using Spectrogram & Phoneme Embedding. In *Proceedings of Interspeech 2018* (pp. 3688–3692).
- Yoon, S., Byun, S., Dey, S., & Jung, K. (2019). Speech emotion recognition using multi-hop attention mechanism. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 2822–2826).
- Yoon, S., Byun, S., & Jung, K. (2018). Multimodal speech emotion recognition using audio and text. In *2018 IEEE Spoken Language Technology Workshop (SLT)* (pp. 112–118).
- Yu, W., Xu, H., Yuan, Z., & Wu, J. (2021). Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 35-12, pp. 10790–10797).
- Yumoto, E., Gould, W., & Baer, T. (1982). Harmonics-to-noise ratio as an index of the degree of hoarseness. *The Journal of the Acoustical Society of America*, 71(6), 1544–1550.
- Yus, F. (2016). *Humour and relevance*. Amsterdam: John Benjamins Publishing Company.
- Zadeh, A., Chen, M., Poria, S., Cambria, E., & Morency, L.-P. (2017). Tensor Fusion Network for Multimodal Sentiment Analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 1103–1114).
- Zadeh, A., Liang, P. P., Mazumder, N., Poria, S., Cambria, E., & Morency, L.-P. (2018). Memory fusion network for multi-view sequential learning. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 32-1).
- Zadeh, A., Liang, P. P., Poria, S., Cambria, E., & Morency, L.-P. (2018). 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)* (pp. 2236–2246).
- Zadeh, A., Liang, P. P., Poria, S., Vij, P., Cambria, E., & Morency, L.-P. (2018). Multi-attention recurrent network for human communication comprehension. In *Thirty-second aaai conference on artificial intelligence*.
- Zadeh, A., Zellers, R., Pincus, E., & Morency, L.-P. (2016). Multimodal sentiment intensity analysis in videos: Facial gestures and verbal messages. *IEEE Intelligent Systems*, 31(6), 82–88.

- Zeng, J., Liu, T., & Zhou, J. (2022). Tag-assisted multimodal sentiment analysis under uncertain missing modalities. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 1545–1554).
- Zhang, B., Cao, S., Zhang, X., Zhang, Y., Ma, L., & Shinozaki, T. (2022). CenSER: Curriculum semi-supervised learning for speech recognition based on self-supervised pre-training. In *Proceedings of Interspeech 2022*.
- Zhang, S., Chadwick, M., Ramos, A. G. C., & Bhattacharya, S. (2022). Cross-attention is all you need: Real-time streaming transformers for personalised speech enhancement. *arXiv preprint arXiv:2211.04346*.
- Zhang, S., Yi, J., Tian, Z., Bai, Y., Tao, J., & Liu, X. (2021). End-to-End Spelling Correction Conditioned on Acoustic Feature for Code-Switching Speech Recognition. In *Proceedings of Interspeech 2021*.
- Zhang, Z., Coutinho, E., Deng, J., & Schuller, B. (2014). Cooperative learning and its application to emotion recognition from speech. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(1), 115–126.
- Zhang, Z., Deng, J., & Schuller, B. (2013). Co-training succeeds in computational paralinguistics. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 8505–8509).
- Zhang, Z., Han, J., Deng, J., Xu, X., Ringeval, F., & Schuller, B. (2018). Leveraging unlabeled data for emotion recognition with enhanced collaborative semi-supervised learning. *IEEE Access*, 6, 22196–22209.
- Zhang, Z., Peng, L., Pang, T., Han, J., Zhao, H., & Schuller, B. W. (2024). Refashioning emotion recognition modelling: The advent of generalised large models. *IEEE Transactions on Computational Social Systems*.
- Zhang, Z., & Schuller, B. (2012). Semi-supervised learning helps in sound event classification. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 333–336).
- Zhao, J., Li, R., & Jin, Q. (2021). Missing modality imagination network for emotion recognition with uncertain missing modalities. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (pp. 2608–2618).
- Zheng, L., Li, Q., Ban, H., & Liu, S. (2018). Speech emotion recognition based on convolution neural network combined with random forest. In *2018 Chinese Control and Decision Conference (CCDC)* (pp. 4143–4147).

- Zhou, S., & Beigi, H. (2020). A transfer learning method for speech emotion recognition from automatic speech recognition. *arXiv preprint arXiv:2008.02863*.
- Zhou, S., & Chellappa, R. (1988). Computation of optical flow using a neural network. *IEEE Transactions on Neural Networks*.
- Zhu, D., Chen, J., Shen, X., Li, X., & Elhoseiny, M. (2024). MINIGPT-4: Enhancing vision-language understanding with advanced large language models. In *12th International Conference on Learning Representations (ICLR) 2024*.
- Zhu, H., Gao, D., Cheng, G., Povey, D., Zhang, P., & Yan, Y. (2023). Alternative pseudo-labeling for semi-supervised automatic speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31, 3320–3330.
- Zhu, X. J. (2005). Semi-supervised learning literature survey. *University of Wisconsin-Madison Department of Computer Sciences*.
- Zhu, Z., & Sato, Y. (2021). Speech emotion recognition using semi-supervised learning with efficient labeling strategies. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)* (pp. 358–365).