



Mitigating Subgroup Disparities in Multi-Label Speech Emotion Recognition: A Pseudo-Labeling and Unsupervised Learning Approach

Yi-Cheng Lin¹, Huang-Cheng Chou², Hung-yi Lee¹

¹Graduate Institute of Communication Engineering, National Taiwan University, Taiwan

²Independent Researcher, Taiwan

{f12942075, hungyilee}@ntu.edu.tw, huangchengchou@gmail.com

Abstract

While subgroup disparities and performance bias are increasingly studied in computational research, fairness in categorical *Speech Emotion Recognition* (SER) remains underexplored. Existing methods often rely on explicit demographic labels, which are difficult to obtain due to privacy concerns. To address this limitation, we introduce an Implicit Demography Inference (IDI) module that leverages pseudo-labeling from a pre-trained model and unsupervised learning using k-means clustering to mitigate bias in SER. Our experiments show that pseudo-labeling IDI reduces subgroup disparities, improving fairness metrics by over 28% with less than a 2% decrease in SER accuracy. Also, the unsupervised IDI yields more than a 4.6% improvement in fairness metrics with a drop of less than 3.6% in SER performance. Further analyses reveal that the unsupervised IDI consistently mitigates race and age disparities, demonstrating its potential when explicit demographic information is unavailable.

Index Terms: Subgroup Disparities, Speech Emotion Recognition, Fairness and Bias

1. Introduction

Speech Emotion Recognition (SER) is essential in human-computer interaction, enhancing user experiences across various applications [1]. As SER becomes increasingly integral to applications ranging from virtual assistants to therapeutic tools, ensuring that these systems operate fairly and without bias is more critical than ever [2, 3]. However, conventional debiasing methods predominantly depend on explicit demographic annotations (e.g., gender, age, race), which are not only challenging to collect due to privacy concerns and ethical issues but are also notably scarce in real-world data [4, 5, 6, 7, 8, 9, 10].

Existing gender debiasing methods in SER are scarce, particularly those that do not require gender labels. Additionally, prior approaches have been designed for single-label categorical SER, while recent research highlights the importance of multi-label SER [11, 12], aligning with psychological studies on emotion co-occurrence [13]. Against these backdrops, our research question is: *How can we mitigate subgroup disparities in multi-label SER without relying on explicit demographic annotations?*

To address this challenge, we introduce the concept of *Implicit Demography Inference (IDI)*. This novel framework aims to simulate demographic supervision by inferring latent group cues directly from the speech data. The IDI module operates through two techniques: first, by leveraging pseudo-labeling to generate proxy demographic labels using a pre-trained model, and second, by employing unsupervised clustering—specifically, applying K-Means to ECAPA-TDNN em-

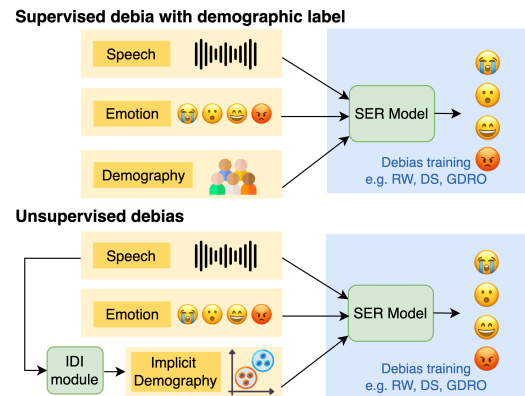


Figure 1: Illustration of the proposed debiasing framework.

beddings [14]—to uncover inherent group structures. Together, these strategies enable the system to approximate demographic distinctions without requiring manual annotations.

Our experiments on the CREMA-D database [15] provide compelling evidence for the efficacy of this approach. We observe that pseudo-labeling yields improvements in key fairness metrics by over 28%, while the unsupervised clustering approach enhances these metrics by approximately 4.6%, all with only a minor compromise in overall SER accuracy. Furthermore, the SER model debiased using the unsupervised clustering method can also suppress age and race bias, proving the generalizability of our method. These results underscore the potential of the IDI module to serve as a robust and scalable solution for debiasing in SER, especially in environments where explicit demographic data is limited.

2. Related Works

2.1. Debias by Supervised Learning

Supervised learning approaches to mitigating gender bias in SER typically rely on explicitly annotated gender labels during training. Previous studies have implemented fairness constraints or modified training strategies to reduce performance disparities between gender subgroups. Chien et al. [7] employed fairness-constrained adversarial training to develop gender-fair representations, considering both speaker and rater gender to enhance the flexibility of emotion predictions. Similarly, Chien et al. [8] addressed biases originating from speakers and raters by integrating fairness-aware learning techniques that combine adversarial training with multi-objective optimization, producing gender-neutral emotion predictions. More recently, Upadhyay et al. [16] demonstrated that while SER models may exhibit fairness within their source datasets, significant gender biases can emerge in cross-corpus scenarios. To miti-

gate this issue, they proposed the Combined Fairness Adaptation method, which leverages adversarial and contrastive learning to align gender representations across domains, enhancing fairness without compromising recognition performance. However, these studies rely on gender labels, often difficult to obtain due to privacy concerns and data limitations. Additionally, they frequently simplify test sets by selecting or removing data to focus on a subset of emotions. Furthermore, these approaches treat SER as a single-label classification task, disregarding samples without consensus emotion labels in the test set.

In contrast to these prior studies, our approach retains all data and labels in the test set to provide a more accurate assessment of model performance. Moreover, we define SER as a multi-label task to better capture the complexity of emotion perception, as supported by recent SER research [11, 12, 17] and psychological insights [13]. Emotion perception is inherently multifaceted, with speakers frequently conveying multiple emotions simultaneously. For example, an utterance may express both anger and sarcasm. Thus, multi-label classification presents a more suitable paradigm for SER than traditional single-label classification.

2.2. Debias by Unsupervised Learning

Chou et al. [18] present a method to mitigate individual speaker biases by employing unsupervised clustering on utterance embeddings from a pre-trained speaker verification model. This method effectively clusters utterances with similar characteristics, approximating true speaker identities without explicit identification. The study shows that incorporating these cluster identifiers enables the creation of a fairness-aware SER model that operates equitably across individual speakers. However, this work focuses on dimensional SER—such as arousal, valence, and dominance—and uses a small subset of the original emotion dataset, which, as Ferrer et al. [19] suggest, may lead to overestimated performance in testing. Therefore, we evaluate all experiments using the original test set.

3. Methodology

We propose a novel framework to mitigate subgroup disparities in gender in multi-label SER without relying on explicit demographic annotations. Central to our approach is the *Implicit Demography Inference (IDI) Module*, which infers group labels from speech data by combining pseudo-labeling and unsupervised clustering. These inferred group labels are then utilized to guide a debias training process. An overview of the proposed framework is illustrated in Fig. 1. In the following sections, we detail the architecture of our proposed IDI module and its integration with debias training techniques.

3.1. Implicit Demography Inference (IDI) Module

The IDI Module is designed to automatically infer latent group labels from speech features through one of two methods: pseudo-labeling and unsupervised clustering.

Pseudo-labeling: we leverage a pre-trained gender detection model [20]¹ to generate pseudo gender labels for each input utterance. These pseudo labels act as initial proxies, capturing potential demographic cues embedded in the speech signal. The pseudo-labels offer a straightforward means to incorporate demographic information without manual annotation. The model’s reliability is assessed on the CREMA-D dataset (detailed in Sec. 4.2), achieving an accuracy of 94.4%.

¹ <https://huggingface.co/audcoring/wav2vec2-large-robust-24-ft-age-gender>

Unsupervised Clustering: We extract embeddings using the ECAPA-TDNN model [14] and apply K-Means clustering [21] to these embeddings. The cluster assignments derived from this process are treated as group labels, reflecting latent structures within the data that may correspond to demographic differences. We set the cluster sizes to 2, 4, 8, 16, and 32.

3.2. Debias Training with Inferred Group Labels

Once the IDI Module infers the group labels, they are incorporated into the debias training of the SER model. Specifically, the inferred labels partition the training data into subgroups, allowing us to impose fairness constraints that decrease performance disparities across these groups. We explore four debiasing methods that leverage these inferred group labels:

Reweighting (RW) [22]: This approach adjusts the importance of each sample based on the frequency of its subgroup within each emotion class. The overall weighted loss is computed as:

$$\mathcal{L}_{RW} = \sum_c \sum_g \frac{n_c}{n_{c,g}} \mathcal{L}_{SER}^{(c,g)}, \quad (1)$$

where $n_{c,g}$ is the number of samples in class c belonging to subgroup g , reflects the inverse frequency of subgroup g in the specific emotion class. n_c is the number of samples in class c . $\mathcal{L}_{SER}^{(c,g)}$ denotes the SER loss in class c and subgroup g .

Downsampling (DS) [22]: This method balances subgroup distributions by equalizing the number of samples per group within each class. For each class, we determine the minimum number of utterances available among all groups and then randomly select that many samples from each group. This ensures that every subgroup is equally represented during training.

Group Distribution Robust Optimization (GDRO) [23]: GDRO seeks to minimize the worst-case loss among all subgroups, thereby ensuring robust performance across groups. The GDRO objective is formulated as:

$$\mathcal{L}_{GDRO} = \max_g \mathcal{L}_{SER}^{(g)}, \quad (2)$$

which directly targets the subgroup with the highest loss.

Group Aware DRO (GADRO) [24]: Extending GDRO, GADRO incorporates an additional regularization term $\frac{\lambda_{GD}}{\sqrt{n_g}}$ to penalize large discrepancies between subgroup losses, further promoting fairness. The loss function for GADRO is defined as:

$$\mathcal{L}_{GADRO} = \max_g (\mathcal{L}_{SER}^{(g)} + \frac{\lambda_{GD}}{\sqrt{n_g}}), \quad (3)$$

where λ_{GD} is a hyperparameter and n_g denotes the group size. In our experiments, we set λ_{GD} to 4.

4. Experiments Setup

We select the SSL-based SER model with WavLM base plus² [25] feature extractor and two linear layers as our primary backbone for the following experiments, because it ranks second on the SER leaderboard [12] while maintaining a lower parameter count compared to the top-1 model, XLSR-1B. We use the class-balanced cross-entropy loss [26] as base SER loss \mathcal{L}_{SER} .

4.1. Baseline

We compare our proposed methods against several baseline models. **Random Baseline:** A naive model that assigns predictions uniformly at random, establishing a lower bound on

² <https://huggingface.co/s3prl/converted.ckpt/resolve/main/wavlm.base.plus.pt>

SER performance. **Empirical Risk Minimization (ERM)**: The standard SER model minimizes the overall \mathcal{L}_{SER} on the training data without any fairness constraints. ERM serves as our conventional reference model. **Existing Unsupervised Debiasing Methods**: We implement state-of-the-art debiasing approaches from computer vision that do not require explicit demographic labels, including LfF [27], and DisEnt [28]. **PGDRO**: A method that leverages probabilistic pseudo-labels for demography inference [29].

4.2. Dataset

We evaluate our method on the widely used CREMA-D database [15], which contains 7,442 utterances (from 91 actors) in English across six emotions (anger, disgust, fear, happy, neutral, sad). Each utterance is annotated by at least 7 raters and includes demographic metadata (gender, age, race, and ethnicity). For reproducibility, we use the audio-only annotations and adhere to the EMO-SUPERB splits [12].

Most prior works on debiasing in SER have filtered or selectively sampled data for their experiments. Inspired by Ferrer et al. [19], we simulate biased training data while evaluating models on the original test set to better reflect real-world conditions in an unseen setting. To achieve this, we employ ad-hoc data pre-processing to manipulate the training data. First, we assign each sample to a single emotion category based on majority voting. Then, we follow the natural demographic trends in the CREMA-D dataset. In this study, we set the gender imbalance ratio to approximately 1:20 (female to male or male to female).

4.3. Evaluation Metrics

For assessing SER, we utilize two primary performance metrics: the macro-F1 score (**F1**) and Hamming accuracy (**ACC**). The macro-F1 score is computed by individually evaluating the F1 measure for each emotion category and then averaging the results, thereby handling class imbalances. Hamming accuracy, on the other hand, calculates the ratio of correctly predicted labels to the total number of labels, offering a comprehensive view of overall model performance. By the approach outlined in [12], we binarize the prediction probabilities using a threshold of $1/|\mathcal{Y}|$, where \mathcal{Y} denotes the set of all emotion classes.

In terms of fairness, our evaluation relies on two critical criteria: Equal Opportunity and Demographic Parity. The Equal Opportunity metric is calculated as the root mean square (RMS) of the gap in the true positive rate (TPR) across all emotion classes [30, 31]. This measure ensures that, given the actual positive instances, the likelihood of correctly predicting an emotion does not vary with protected attributes such as gender or race. In addition, Demographic Parity requires that the probability of a positive prediction remains consistent across different demographic groups, independent of the true label. We express this disparity mathematically as:

$$DP_{gap} = \sqrt{\sum_{\hat{y} \in \mathcal{Y}} \max_{g \in \mathcal{G}} DP(g, \hat{y})^2}; \quad (4)$$

$$DP(g, \hat{y}) = \mathbb{E} \left[P(\hat{y} = 1 \mid G = g) - P(\hat{y} = 1) \right], \quad (5)$$

where G represents a group, \mathcal{G} represents the set of all groups, and \hat{y} represents the emotion prediction. A higher value of DP_{gap} implies a greater imbalance.

Table 1: Summary of results in SER accuracy (**F1** and **ACC**) and fairness (**TPR_{gap}** and **DP_{gap}**) on gender. **Super.** specifies the level of supervision for debias. ● indicates debias with ground truth gender labels. ⦿ represents the pseudo-labeling approach, and ○ refers to bias unsupervised methods. For unsupervised IDI, we use k -means with $k=16$. **Overall** averages the four IDI methods. The **Gain** represents the improvement compared to the **ERM** method. ↑ indicates that higher values correspond to better performance; ↓ indicates the opposite.

Method	Super.	F1↑	ACC↑	TPR _{gap} ↓	DP _{gap} ↓
ERM		0.651	0.824	0.278	0.103
Random		0.360	0.509	0.060	0.015
LfF [27]	○	0.655	0.825	0.268	0.099
DisEnt [28]	○	0.653	0.825	0.261	0.094
PGDRO [29]	⦿	0.643	0.817	0.318	0.121
DS [22]	●	0.598	0.804	0.094	0.033
DS-IDI	⦿	0.580	0.811	0.137	0.045
	○	0.561	0.755	0.201	0.100
RW [22]	●	0.646	0.823	0.171	0.061
RW-IDI	⦿	0.647	0.826	0.197	0.071
	○	0.642	0.817	0.237	0.091
GDRO [23]	●	0.627	0.796	0.199	0.086
GDRO-IDI	⦿	0.632	0.807	0.229	0.088
	○	0.639	0.803	0.228	0.104
GADRO [24]	●	0.621	0.788	0.185	0.084
GADRO-IDI	⦿	0.631	0.807	0.229	0.088
	○	0.638	0.804	0.223	0.098
Overall	●	0.623	0.803	0.162	0.066
	⦿	0.623	0.813	0.198	0.073
	○	0.620	0.795	0.222	0.098
Gain ↑	●	-4.30%	-2.58%	41.64%	35.92%
	⦿	-4.38%	-1.37%	28.78%	29.13%
	○	-4.76%	-3.55%	20.05%	4.61%

5. Results and Analyses

Table 1 presents the averaged results on performance and gender bias across five folds. We evaluate the effectiveness of our proposed methods and provide detailed analyses in the following sections.

5.1. Performance and Effects of Pseudo-labeling

Our results show that pseudo-labeling reduces subgroup disparities in multi-label SER. As seen in Table 1, using pseudo-labels from a pre-trained gender detection model improves the TPR_{gap} fairness metric by 28.78% and DP by 29.13% relative to the ERM baseline, with less than a 1.5% drop in overall accuracy. Compared to PGDRO [29], our modified reweighting strategy with pseudo-labels achieves a better balance between fairness and performance.

5.2. Performance and Effects of Unsupervised Learning

The number of clusters plays a critical role in the effectiveness of this method. Using too few clusters may result in overly coarse groupings that fail to capture the underlying nuances in the data, while employing too many clusters can fragment the data and introduce noise, leading to less meaningful subgroup distinctions. We found that a cluster size of 16 provides a balanced granularity, offering significant improvements in fairness metrics (20.05% in TPR_{gap} and 4.61% in DP_{gap}) while maintaining a minor drop in overall performance (a 4.76% decrease

Table 2: The results of unsupervised clustering on **Race** and **Age**. **G** denotes the number of subgroups within the demographic information. **K** refers to the number of clusters. **Gain** represents the performance improvement compared to the baseline model, **ERM**.

Model	K	Race (G=3)		Age (G=3)			
		F1↑	ACC↑	TPR _{gap}	DP _{gap}	TPR _{gap}	DP _{gap}
ERM		0.651	0.824	0.183	0.111	0.153	0.090
DS	16	0.561	0.755	0.142	0.119	0.134	0.095
	32	0.575	0.773	0.165	0.112	0.137	0.106
RW	16	0.642	0.817	0.147	0.100	0.140	0.079
	32	0.628	0.810	0.172	0.112	0.150	0.095
GADRO	16	0.639	0.803	0.159	0.118	0.138	0.095
	32	0.640	0.806	0.162	0.109	0.137	0.096
GDRO	16	0.638	0.804	0.162	0.113	0.122	0.089
	32	0.629	0.800	0.172	0.111	0.126	0.090
Gain↑	16	-4.76%	-3.55%	16.67%	-1.35%	12.75%	0.56%
	32	-5.07%	-3.25%	8.33%	0.00%	10.13%	-7.50%

in F1 and a 3.55% decrease in ACC), as shown in Table 1.

Compared to previously proposed unsupervised debiasing methods such as LfF and DisEnt, our approach demonstrates a more favorable trade-off between fairness and SER performance. While these earlier methods yield minor performance improvements, they often incur few debias effects in TPR_{gap} and DP_{gap}. In contrast, our method not only enhances fairness metrics significantly but also maintains higher recognition performance.

To further investigate the underlying reasons behind the method’s effectiveness, we visualize the ECAPA-TDNN embeddings using t-SNE [32], as shown in Fig. 4. The male and female samples are clearly separated, confirming that our clustering captures gender information and supports the scalability and label-free nature of our unsupervised approach.

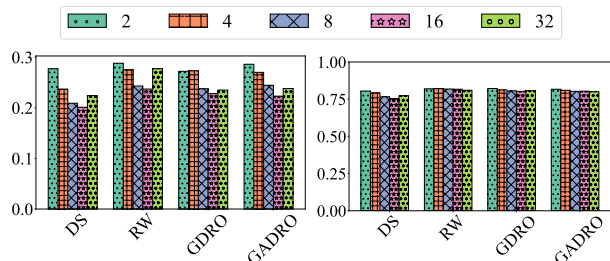


Figure 2: The left subfigure presents the TPR_{gap}, while the right one depicts the Hamming accuracy. The bars illustrate the performance of the proposed unsupervised debiasing method, evaluated across four different techniques with varying cluster sizes.




Figure 3: Illustration of the averaged results across four models trained with the proposed unsupervised method, using different cluster sizes, in terms of Accuracy and TPR.




Figure 4: Illustration of samples in Fold 5 using t-SNE, with ECAPA-TDNN embeddings and a cluster size of 16. Circles and triangles represent male and female speakers, respectively, while different colors indicate different clusters.

5.3. Potential of Unsupervised Clustering

We assess the applicability of unsupervised clustering to additional demographic attributes, as it does not rely on explicit demographic information. The CREMA-D dataset provides demographic metadata, including speakers’ age and race (specifically, Caucasian, African American, and Asian). In accordance with [33], ages are classified into three subgroups: “Young” (20–35 years), “Middle” (36–49 years), and “Elderly” (50–75 years).

Table 2 presents overall performance across subgroups, demonstrating that the unsupervised method with K=16 clusters improves TPR_{gap} metric by over 12.75%, while maintaining an SER performance decrease of 4.76% and 3.55% in F1 and ACC, respectively. This suggests that the method’s performance is robust across different demographic attributes, highlighting its potential for broad applicability for subgroup disparities in multi-label SER tasks.

6. Limitations

While our approach demonstrates robust improvements in mitigating subgroup disparities, several limitations warrant discussion. First, our experiments are based on the CREMA-D dataset, which comprises acted emotional expressions; thus, the generalizability to naturalistic settings remains to be validated. Second, our methods depend heavily on the quality of pre-trained models—the gender detection model for pseudo-labeling and the ECAPA-TDNN for embedding extraction. More pre-trained models can be explored.

7. Conclusion and Future Work

This work empirically examines various debiasing methods in supervised, pseudo-label, and unsupervised learning to mitigate subgroup disparities without explicit demographic information. We leverage pseudo-labels generated by a pre-trained model and introduce an unsupervised clustering approach that integrates with existing debiasing techniques originally designed for demographic labels. Our findings demonstrate that pseudo-label learning improves TPR and DP metrics by 28.78% and 29.13%, respectively, with only a minor 4.38% degradation in SER accuracy. Similarly, our proposed unsupervised method achieves 20.05% and 4.61% improvements in TPR and DP metrics, respectively, with just a 4.76% drop in SER performance. For future work, we will explore more possible bias attributes like content [34], language [35], or speaking style [36].

8. References

- [1] D. Thiripurasundari, K. Bhangale, V. Aashritha, S. Mondreti, and M. Kothandaraman, "Speech emotion recognition for human-computer interaction," *Int. J. Speech Technol.*, vol. 27, no. 3, pp. 817–830, Aug. 2024.
- [2] D. Varona and J. L. Suárez, "Discrimination, Bias, Fairness, and Trustworthy AI," *Applied Sciences*, vol. 12, no. 12, 2022.
- [3] C. Ashurst and A. Weller, "Fairness Without Demographic Data: A Survey of Approaches," in *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, ser. EAAMO '23. New York, NY, USA: Association for Computing Machinery, 2023.
- [4] Y.-C. Lin, H. Wu, H.-C. Chou, C.-C. Lee, and H. yi Lee, "Emobias: A Large Scale Evaluation of Social Bias on Speech Emotion Recognition," in *Interspeech 2024*, 2024, pp. 4633–4637.
- [5] C. Gorrostieta, R. Lotfian, K. Taylor, R. Brutti, and J. Kane, "Gender De-Biasing in Speech Emotion Recognition," in *Interspeech 2019*, 2019, pp. 2823–2827.
- [6] Y.-C. Lin, T.-Q. Lin, H.-C. Lin, A. T. Liu, and H. yi Lee, "On the social bias of speech self-supervised models," in *Interspeech 2024*, 2024, pp. 4638–4642.
- [7] W.-S. Chien and C.-C. Lee, "Achieving Fair Speech Emotion Recognition via Perceptual Fairness," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [8] W.-S. Chien, S. G. Upadhyay, and C.-C. Lee, "Balancing Speaker-Rater Fairness for Gender-Neutral Speech Emotion Recognition," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 11 861–11 865.
- [9] A. Koudounas, E. Pastor, G. Attanasio, L. de Alfaro, and E. Baralis, "Prioritizing data acquisition for end-to-end speech model improvement," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 7000–7004.
- [10] Y. Zhang, Y. Zhang, B. Halpern, T. Patel, and O. Scharenborg, "Mitigating bias against non-native accents," in *Interspeech 2022*, 2022, pp. 3168–3172.
- [11] H.-C. Chou, H. Wu, L. Goncalves, S.-G. Leem, A. Salman, C. Busso, H.-Y. Lee, and C.-C. Lee, "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)*, 2024, pp. 502–509.
- [12] H. Wu *et al.*, "Open-Emotion: A Reproducible EMO-Superb For Speech Emotion Recognition Systems," in *2024 IEEE Spoken Language Technology Workshop (SLT)*, 2024, pp. 510–517.
- [13] A. S. Cowen and D. Keltner, "Semantic Space Theory: A Computational Approach to Emotion," *Trends in Cognitive Sciences*, vol. 25, no. 2, pp. 124–136, 2021.
- [14] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification," in *Interspeech 2020*, 2020, pp. 3830–3834.
- [15] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "CREMA-D: Crowd-Sourced Emotional Multimodal Actors Dataset," *IEEE Transactions on Affective Computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [16] S. G. Upadhyay, W.-S. Chien, and C.-C. Lee, "Is It Still Fair? Investigating Gender Fairness in Cross-Corpus Speech Emotion Recognition," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025.
- [17] S. Park, B. Jeon, S. Lee, and J. Yoon, "Multi-Label Emotion Recognition of Korean Speech Data Using Deep Fusion Models," *Applied Sciences*, vol. 14, no. 17, 2024.
- [18] H.-C. Chou, H. Wu, and C.-C. Lee, "Stimulus Modality Matters: Impact of Perceptual Evaluations from Different Modalities on Speech Emotion Recognition System Performance," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025.
- [19] L. Ferrer, O. Scharenborg, and T. Bäckström, "Good practices for evaluation of machine learning systems," 2024.
- [20] F. Burkhardt, J. Wagner, H. Wierstorf, F. Eyben, and B. Schuller, "Speech-based Age and Gender Prediction with Transformers," in *Speech Communication; 15th ITG Conference*, 2023, pp. 46–50.
- [21] S. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982.
- [22] F. Kamiran and T. Calders, "Data preprocessing techniques for classification without discrimination," *Knowledge and Information Systems*, vol. 33, no. 1, pp. 1–33, 2012.
- [23] A. Ben-Tal, D. den Hertog, A. De Waegenare, B. Melenberg, and G. Rennen, "Robust Solutions of Optimization Problems Affected by Uncertain Probabilities," *Management Science*, vol. 59, no. 2, pp. 341–357, 2013.
- [24] S. Sagawa*, P. W. Koh*, T. B. Hashimoto, and P. Liang, "Distributionally Robust Neural Networks for Group Shifts: On the Importance of Regularization for Worst-Case Generalization," in *International Conference on Learning Representations*, 2020.
- [25] S. Chen *et al.*, "WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing," *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [26] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, "Class-Balanced Loss Based on Effective Number of Samples," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [27] J. Nam, H. Cha, S. Ahn, J. Lee, and J. Shin, "Learning from failure: training debiased classifier from biased classifier," in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2020.
- [28] J. Lee, E. Kim, J. Lee, J. Lee, and J. Choo, "Learning Debiased Representation via Disentangled Feature Augmentation," in *Advances in Neural Information Processing Systems*, 2021.
- [29] S. S. Ghosal and Y. Li, "Distributionally Robust Optimization with Probabilistic Group," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 10, pp. 11 809–11 817, Jun. 2023.
- [30] X. Han, T. Baldwin, and T. Cohn, "Diverse adversaries for mitigating bias in training," in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 2021.
- [31] M. De-Arteaga, A. Romanov, H. Wallach, J. Chayes, C. Borgs, A. Chouldechova, S. Geyik, K. Kenthapadi, and A. T. Kalai, "Bias in bios: A case study of semantic representation bias in a high-stakes setting," in *Proceedings of the Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, 2019.
- [32] L. van der Maaten and G. Hinton, "Visualizing Data using t-SNE," *Journal of Machine Learning Research*, vol. 9, no. 86, pp. 2579–2605, 2008.
- [33] L. Goncalves, H.-C. Chou, A. N. Salman, C.-C. Lee, and C. Busso, "Jointly Learning From Unimodal and Multimodal-Rated Labels in Audio-Visual Emotion Recognition," *IEEE Open Journal of Signal Processing*, vol. 6, pp. 165–174, 2025.
- [34] Y.-C. Lin *et al.*, "Listen and speak fairly: a study on semantic gender bias in speech integrated large language models," in *2024 IEEE Spoken Language Technology Workshop (SLT)*, 2024.
- [35] H.-C. Lin, Y.-C. Lin, H.-C. Chou, and H. yi Lee, "Improving speech emotion recognition in under-resourced languages via speech-to-speech translation with bootstrapping data selection," 2025. [Online]. Available: <https://arxiv.org/abs/2409.10985>
- [36] Y. Meng, Y.-H. Chou, A. T. Liu, and H.-y. Lee, "Don't speak too fast: The impact of data bias on self-supervised speech models," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.