



Meta-PerSER: Few-Shot Listener Personalized Speech Emotion Recognition via Meta-learning

Shi-Xin Fang^{*1}, Liang-Yeh Shen^{*1}, Yi-Cheng Lin¹, Huang-Cheng Chou², Hung-yi Lee¹

¹National Taiwan University, Taiwan,

²Independent Researcher, Taiwan

{b10901005, b10507008, f12942075, hungyilee}@ntu.edu.tw, huangchengchou@gmail.com

Abstract

This paper introduces **Meta-PerSER**, a novel meta-learning framework that personalizes Speech Emotion Recognition (SER) by adapting to each listener’s unique way of interpreting emotion. Conventional SER systems rely on aggregated annotations, often overlooking individual subtleties and leading to inconsistent predictions. In contrast, Meta-PerSER leverages a Model-Agnostic Meta-Learning (MAML) approach enhanced with Combined-Set Meta-Training, Derivative Annealing, and per-layer per-step learning rates, enabling rapid adaptation with only a few labeled examples. By integrating robust representations from pre-trained self-supervised models, our framework first captures general emotional cues and then fine-tunes itself to personal annotation styles. Experiments on the IEMOCAP corpus demonstrate that Meta-PerSER significantly outperforms baseline methods in both seen and unseen data scenarios, highlighting its promise for personalized emotion recognition.

Index Terms: speech emotion recognition, few-shot learning, meta-learning, AI personalization, multi-label classification

1. Introduction

Speech Emotion Recognition (SER) aims to identify emotional states from vocal cues automatically [1, 2, 3]. Despite substantial progress with deep learning, SER remains challenging due to the complexity of emotional expressions and the subjectivity of emotion perception [4]. One fundamental challenge in SER is the variability across speakers and listeners [5, 6].

Traditional SER models typically assume a one-size-fits-all approach, learning a global mapping from acoustic features to emotion labels. However, this approach neglects individual differences: speakers have unique vocal characteristics, and annotators (listeners) often interpret the same speech differently, leading to inconsistent labels and noisy emotion datasets. While speaker normalization or adaptation techniques [7, 8] address speaker variability, they still strive for a universal model that might not account for each listener’s perception of emotion. In practice, a model that could adapt to each individual’s emotion perception—essentially a listener-personalized SER system—would be highly desirable, yet existing SER approaches lack this personalization capability.

A key challenge in listener-personalized SER is the high variability in emotion annotations across individuals. Aggregating annotations from multiple listeners into a single ground-truth label risks losing personal perception and introducing biases [9, 10, 11, 12]. Conversely, training separate models for each listener requires substantial labeled data and computational resources, which is impractical for real-world deployment

[13, 14]. Additionally, real-world applications require quick adaptation to new, unseen listeners with minimal labeled data, a scenario where traditional SER approaches struggle [15].

Meta-learning offers a promising avenue to tackle the aforementioned challenges by enabling models to learn from a distribution of tasks [16, 17]. Such algorithms can rapidly adapt to new tasks with only a few training examples, making them well-suited for data-scarce scenarios. In SER, each task could correspond to a specific condition (e.g., a new speaker or a listener’s annotation style). By training on many such tasks, the model learns how to adapt its parameters to a new task quickly. Recent studies have started exploring meta-learning for SER to improve cross-corpus performance or reduce annotation noise [18, 19, 20, 21, 22], but to our knowledge, none have focused on personalizing the model to individual listeners’ interpretations. While the study [23] builds personalized SER models, it focuses on dimensional SER (e.g., arousal and valence), whereas our work focuses on categorical SER (e.g., happy or angry).

In this paper, we propose *Meta-PerSER*, a novel meta-learning framework for SER that incorporates the listener’s perspective into the learning process to achieve listener-personalized adaptation. During meta-training, the framework simulates adapting to different listeners by treating each listener’s labeled data as a separate task. This approach allows the model to learn a good initialization that can be efficiently fine-tuned for a new listener using only a small amount of that listener’s data. As a result, our method can quickly calibrate to a new listener’s subjective emotion perception, thus addressing label ambiguity by aligning predictions with each listener’s expectations. Unlike traditional SER models that yield a fixed classifier, Meta-PerSER is inherently adaptable, bridging the gap between generic and personalized emotion recognition. We will release our code¹.

Experiments on a well-known emotion database, the IEMOCAP [24], show that Meta-PerSER outperforms baselines, especially in few-shot settings with limited new-listener data. Meta-PerSER underscores the benefit of personalization in SER. In summary, our main contributions are:

- **Listener-personalized SER via meta-learning:** We introduce the first SER framework that personalizes individual listeners’ categorical emotion interpretations.
- **Efficient adaptation with limited data:** We develop a meta-training strategy that rapidly adapts to a new listener using only a small amount of data.
- **Improved performance and robustness:** We empirically demonstrate that Meta-PerSER improves emotion recognition accuracy and robustness to annotation differences, outperforming state-of-the-art baselines.

^{*}Equal contribution

¹<https://github.com/Jeffabcd/Meta-PerSER>

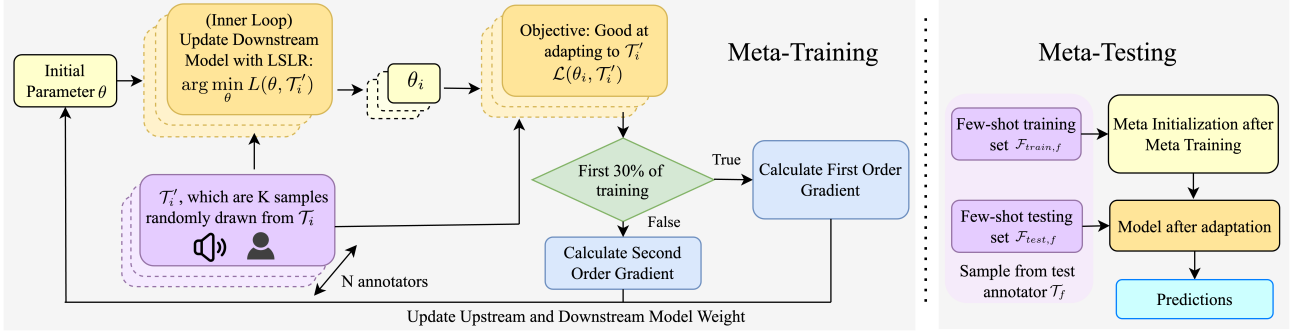


Figure 1: Framework of the proposed Meta-PerSER.

2. Methodology

2.1. Backbone SER Framework

We employ a unified model architecture based on the s3prl toolkit [25]. Our SER framework comprises two principal components: an upstream module and a downstream module. The upstream module leverages pre-trained self-supervised learning (SSL) models, the base variants of Wav2Vec2 [26]², HuBERT [27]³, and WavLM [28]⁴. In terms of the downstream module, we use a mean pooling layer and two linear layers. Additionally, the objective function for the SER task is class-balanced cross-entropy loss [29].

2.2. Problem Formulation

We address the challenge of personalized SER in a few-shot learning setting, as collecting large amounts of personalized labeled data for every user is impractical in real-world scenarios. For each experiment, we divide the data into three sets: a training set, a few-shot training set, and a few-shot testing set. The training set $\mathcal{T} = \{\mathcal{T}_i\}_{i=1}^N$ consists of a large set of labeled examples annotated by N annotators, used for general model training. For each new annotator f , we withhold their data \mathcal{T}_f from the training set. From this held-out data, we construct a few-shot training set $\mathcal{F}_{train,f} = \{(x_{fj}, y_{fj})\}_{j=1}^K$, which contains K labeled speech samples used for model adaptation. Finally, the few-shot testing set $\mathcal{F}_{test,f} = \{(x_{fj}, y_{fj})\}_{j=1}^Q$ consists of an additional Q labeled speech samples from annotator f , used exclusively for evaluation. Here, x_{fj} represents the speech signal and y_{fj} the corresponding emotion labels provided by annotator f . The objective is to rapidly adapt a generic SER model to capture the new annotator’s emotion perception using only a few personalized examples.

2.3. Meta-PerSER Framework

Meta-PerSER is a novel meta-learning framework designed to improve the efficiency and personalization of SER. Based on Model-Agnostic Meta-Learning (MAML), our approach integrates multiple enhancements to improve generalization, reduce computational costs, and optimize training efficiency.

2.3.1. Pretraining for General Emotion Representation

Before personalization, we first train a base SER model to capture general emotional patterns. This model is trained on the training set \mathcal{T} , which consists of labeled emotion data aggregated from multiple annotators. We denote this pretrained

model as **SSL-E**, a strong initialization for subsequent adaptation.

2.3.2. Combined-Set Meta-Training (CSMT)

To enable rapid adaptation to individual annotators’ unique emotion perceptions, we employ meta-learning. Traditional MAML splits data into support and query sets to simulate task adaptation and evaluation separately. However, in our setting, annotators label data from the same domain and share similar feature distributions. As a result, the distinction between adaptation and evaluation is less meaningful, and splitting the data provides little benefit. Instead, we introduce **Combined-Set Meta-Training (CSMT)**, where we sample and update the model using all available training data. This approach allows the model to leverage a larger, more diverse dataset during training, improving generalization across annotators without unnecessary partitioning.

2.3.3. Efficient Gradient Optimization

Computational efficiency is a key challenge in meta-learning, particularly due to the high cost of computing second-order derivatives in standard MAML. To address this, we apply **Derivative Annealing (DA)** [30], which initially uses only first-order gradients during early training and gradually introduces second-order derivatives in later stages. This approach significantly reduces computational overhead while still capturing higher-order interactions when needed.

In addition, we integrate **Learning Per-Layer Per-Step Learning Rates and Gradient Directions (LSLR)** [30], a dynamic learning rate adjustment mechanism that optimizes per-layer learning rates at each update step. While DA improves computational efficiency by mitigating second-order derivative costs, LSLR enhances adaptation stability and reduces the need for extensive hyperparameter tuning.

2.3.4. Meta-PerSER Training and Adaptation Process

The overall Meta-PerSER workflow consists of two phases: meta-training and meta-testing.

Meta-Training Phase: The meta-training process begins with the initialized SSL-E model, which serves as the starting point for adaptation. In the inner loop, for each annotator i , we randomly sample K examples from \mathcal{T}_i as \mathcal{T}_i^t , compute the inner loss $\mathcal{L}(\theta, \mathcal{T}_i^t)$, and update model weights using LSLR to obtain annotator-specific parameters θ_i . Following this, the outer loop aggregates gradients from multiple annotators using $\mathcal{L}(\theta_i, \mathcal{T}_i^t)$, updating shared initialization θ to improve generalization. To optimize computational efficiency, we apply Derivative Annealing (DA), which restricts gradient computation to first-order

²https://dl.fbaipublicfiles.com/fairseq/wav2vec/wav2vec_small.pt

³https://dl.fbaipublicfiles.com/hubert/hubert_base_ls960.pt

⁴https://huggingface.co/s3prl/converted_ckpts/resolve/main/wavlm_base.pt

Table 1: This table summarizes the results for both "Seen Data" and "Unseen Data" scenarios using WavLM-, Wav2vec2-, and HuBERT-based SER models. All values are presented in **percentages (%)**, and the best performance for metrics is highlighted in **bold**.

Scenario	Seen Data									Unseen Data								
	WavLM			Wav2vec2			HuBERT			WavLM			Wav2vec2			HuBERT		
Upstream	maF1	miF1	UA	maF1	miF1	UA	maF1	miF1	UA	maF1	miF1	UA	maF1	miF1	UA	maF1	miF1	UA
Random	14.6	17.4	60.7	14.6	17.4	60.7	14.6	17.4	60.7	14.7	18.0	60.8	14.7	18.0	60.8	14.7	18.0	60.8
Linear-Few	21.7	34.5	64.9	21.9	35.1	65.4	21.3	34.0	64.0	20.4	33.5	62.9	19.9	34.1	63.7	20.0	32.6	60.1
Entire-Few	34.3	48.0	80.7	34.2	47.5	79.8	33.9	47.1	80.0	24.8	39.7	75.9	23.8	39.7	75.0	25.6	42.3	77.0
Entire-Zero	32.5	43.8	77.8	26.6	34.8	63.1	26.9	35.2	63.5	24.7	37.4	74.4	24.6	38.0	74.5	25.0	40.3	74.8
Multi-Few	23.3	29.4	62.7	22.3	28.4	63.5	23.8	30.0	64.8	20.0	33.7	66.5	19.4	33.9	67.9	20.8	33.2	67.0
Entire-Sim.	24.4	35.8	63.5	24.4	37.0	64.1	25.5	37.2	64.8	21.8	33.7	56.1	21.5	33.7	57.1	21.3	33.5	56.3
Meta-PerSER	35.7	50.7	82.8	35.6	48.7	81.3	35.3	51.2	82.4	30.7	47.6	80.9	27.2	44.9	78.7	30.2	46.4	80.0

derivatives in the early stages (first 30% of inner-loop steps) before transitioning to second-order derivatives later for richer adaptation.

Meta-Testing Phase: A new, previously unseen annotator is introduced during the meta-testing phase to assess the model’s adaptability. The model undergoes adaptation using K samples from annotator f , $\mathcal{F}_{train,f}$. Finally, its performance is evaluated on Q samples in $\mathcal{F}_{test,f}$, thereby testing its generalization capability on novel annotator data.

This two-phase process ensures that Meta-PerSER enables rapid adaptation to new annotators while maintaining strong generalization.

2.4. Baselines

Our baselines are named according to a two-stage process: first, how the model initialization is trained before few-shot adaptation, followed by a hyphen (“-”), and second, how the few-shot training set \mathcal{F}_{train} is employed for adaptation.

SSL Approach: We train SSL-based SER models using aggregated labels from training annotators, resulting in the SSL-E model in Sec. 2.3.1. We then adapt them to unseen annotators during meta-testing. We evaluated three strategies: the **Entire-Few** baseline, where the SSL-E model is fine-tuned with few-shot training set \mathcal{F}_{train} during meta-testing; **Linear-Few**, in which only the classifier and weighted sum of SSL features are trained on the entire training set prior to few-shot adaptation; and **Entire-Zero**, where the SSL-E model is evaluated in a zero-shot manner on the few-shot testing set \mathcal{F}_{test} without any further fine-tuning.

Multi-Task Approach (Multi-Few): We model each annotator’s labeling behavior through individual classification tasks. While the upstream model is shared, each annotator has a separate prediction head in the classifier. After training, meta-testing is used to learn a new classifier from random initialization for unseen annotators.

Similarity-based Approach (Entire-Sim.) [31]: We use features extracted by the SSL-E model to classify \mathcal{F}_{test} samples based on their proximity to few-shot training samples in the feature space. For each emotion e , we compute a center c_e by averaging the SSL feature embeddings of all few-shot training samples labeled with that emotion. Then, for each few-shot testing sample x , we calculate the cosine similarity s between its SSL feature embedding $f(x)$ and each emotion center c_e . The similarity is transformed via a softmax function to produce a probability distribution over the emotion labels. The probability that x is assigned to the emotion label e_i is:

$$p(y = e_i | x) = \frac{\exp(s(c_{e_i}, f(x)))}{\sum_{e \in \mathcal{E}} \exp(s(c_e, f(x)))}, \quad (1)$$

where \mathcal{E} is the set of all emotions.

3. Experiments Settings

3.1. Resource

We use the IEMOCAP corpus [24], which contains 10,039 utterances by 10 professional actors. Each utterance is labeled by at least three annotators, who may assign one or more emotion labels from a set of 10 categories—namely, frustrated, angry, sad, disgust, excited, fear, neutral, surprise, happy, and “other.” We exclude the “other” category, defining a 9-class SER task. In total, 12 annotators participated in the labeling process. One annotator was excluded due to insufficient data, leaving 11, comprising six self-reporting actors and five external evaluators. We use only the five external annotators for testing, as they each provided annotations across all five sessions. In each test run, one of 5 external annotators is used for testing, one for validation, and the remaining nine for training.

3.2. Application Scenarios

We propose two distinct application scenarios to validate the proposed methods’ efficacy.

3.2.1. Seen Data

In the Seen Data scenario, the SER system has already encountered the utterances during training but has not seen the labels provided by unseen annotators. To simulate this, we partitioned the IEMOCAP dataset by annotator. We conducted five independent experiments; in each, one of these five annotators was designated for testing, while the remaining annotators were used for training. The final results are averaged across all experiments, assessing whether the model can effectively capture variations in annotation styles for the same audio.

3.2.2. Unseen Data

In the Unseen Data scenario, we use the IEMOCAP dataset’s predefined five-session split to prevent utterance overlap between meta-training and meta-testing. One annotator is chosen for testing, and the rest for training; sessions 1–4 are allocated for training, and session 5 for testing. The final training and test sets were composed based on the annotations provided by the selected annotators within each session. This ensures no duplicate audio samples across splits, though it reduces the dataset size.

3.3. Evaluation

We conduct few-shot meta-testing using data annotated by a single annotator. Each experiment is repeated 10 times with different random seeds to ensure robustness. We randomly select $K = 32$ samples for few-shot training and $Q = 128$ samples for few-shot testing in each iteration. We update the model with

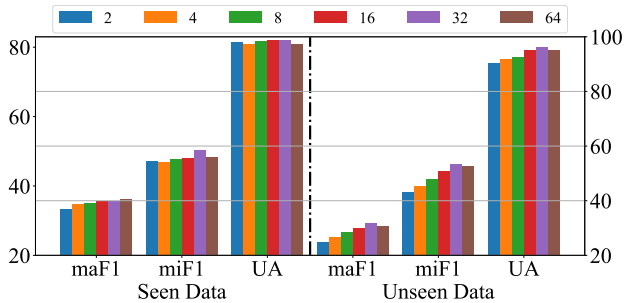


Figure 2: Illustration of the averaged results across different numbers of few-shot training samples (shown in legend) under two scenarios in terms of macro-F1 (**maF1**), micro-F1 (**miF1**), and unweighted accuracy (**UA**) in percentages (%).

one batch and 50 steps.

The SER task is defined as a multi-label classification problem in this study, diverging from conventional approaches, to accommodate the inherent ambiguity of emotion perception. All annotators can select one or more emotions from the provided options. Following the procedure described in [12], we convert the model’s prediction probabilities into binary outcomes by applying a threshold of $1/|\mathcal{Y}|$, where \mathcal{Y} represents the complete set of emotion classes. We employ macro-F1 (**maF1**) and micro-F1 (**miF1**) scores [32, 33] along with unweighted accuracy (UA) as evaluation metrics for assessing and reporting multi-label classification performance.

4. Results and Analysis

4.1. Proposed Meta-PerSER

Table 1 demonstrates that Meta-PerSER consistently outperforms all baseline methods across both Seen and Unseen Data scenarios and across all upstream models. Under the Seen Data setting, Meta-PerSER improves macro-F1, micro-F1, and unweighted accuracy scores by approximately 1–2 percentage points over the best SSL fine-tuning baseline (SSL-FT), indicating enhanced adaptation even when the model has previously encountered the audio. In the more challenging Unseen Data scenario, where the model must generalize to entirely new audio samples, Meta-PerSER shows a larger performance margin, with improvements of up to 7 percentage points in key metrics compared to SSL-based methods.

These results validate that our integration of CSMT, DA, and LSLR effectively enhances the model’s ability to quickly adapt to unseen annotators while preserving robust emotion representations, surpassing both traditional fine-tuning and multi-task strategies.

4.2. Impacts of Few-Shot Training Set Size

Figure 2 illustrates the average performance of Meta-PerSER across different models using various few-shot training set sizes. Overall, performance improves as the few-shot training set size increases from 2 to 32 samples. For seen data, macro-F1 increases from 33.2 to 35.5, micro-F1 reaches a peak of 50.2, and unweighted accuracy (UA) improves to 82.1 at 32 samples. Similarly, for unseen data, all metrics steadily rise up to a few-shot training set size of 32. However, increasing the few-shot training set size to 64 results in a slight decline in performance, particularly for unseen data. This suggests that while a larger few-shot training set initially provides more representative examples for effective adaptation, exceeding an optimal size (in this case, 32 samples) may introduce outlier or ambiguous samples that confuse the model.

Table 2: This table summarizes the performance of adding the proposed components one by one in percentages (%). **INI** indicates whether the model is initialized from the pre-trained **SSL-E**; \checkmark indicates that the model incorporates the corresponding techniques. **CSMT**, **DA**, and **LSLR** are defined in Section 2.

Upstream	WavLM			Wav2vec2			HuBERT		
INI CSMT DA LSLR	maF1	miF1	UA	maF1	miF1	UA	maF1	miF1	UA
	22.8	40.1	76.8	23.6	40.9	77.4	23.6	40.7	78.1
\checkmark	25.7	39.3	74.8	24.9	39.4	72.2	27.1	41.9	76.3
\checkmark \checkmark	27.6	40.7	77.3	26.6	39.1	74.5	28.1	41.5	77.5
\checkmark \checkmark \checkmark	27.9	41.8	77.0	24.9	39.4	72.2	27.0	41.7	76.3
\checkmark \checkmark \checkmark \checkmark	27.2	43.8	77.1	25.3	42.4	75.4	26.8	44.5	78.1
\checkmark \checkmark \checkmark \checkmark \checkmark	30.7	47.6	80.9	27.2	44.9	78.7	30.2	46.4	80.0

4.3. Effectiveness of Different Techniques

Table 2 presents the ablation study on techniques, proving that the four components of Meta-PerSER collectively enhance performance under the Unseen Data scenario. Specifically, the incorporation of **INI** (i.e., leveraging the pre-trained **SSL-E** weights) yields a significant improvement in the macro-F1 metric across all three SSL-based models. Additionally, the individual application of CSMT, DA, and LSLR consistently enhances performance across all evaluation metrics. The combination of **CSMT** and **INI** generally outperforms configurations using either (**DA** and **INI**) or (**LSLR** and **INI**) in most cases. Most importantly, the comprehensive integration of all three methods produces the highest gains in macro-F1, micro-F1, and unweighted accuracy across all upstream models. These findings robustly validate that our integrated approach, **Meta-PerSER**, substantially augments the model in unseen scenarios, thereby advancing the state-of-the-art for personalized SER.

5. Limitations

We are among the first to investigate categorical personalized SER systems. However, our current experimental settings do not incorporate conversational context, as emotional ratings are based solely on conversational-level audio-visual cues that the annotators provide. This limitation may hinder the system’s ability to accurately capture the natural context of individual emotion perception.

6. Conclusion and Future Work

This paper introduces a novel framework, **Meta-PerSER**, designed to effectively adapt to unseen annotators in SER tasks with only a few labeled examples. Meta-PerSER integrates a pre-trained self-supervised backbone with Combined-Set Meta-Training, Derivative Annealing, and per-layer adaptive learning rates. This design enables our system to outperform traditional fine-tuning and multi-task baselines in seen and unseen scenarios on the IEMOCAP corpus. Our results demonstrate that meta-learning can capture subjective emotion perceptions without extensive annotation collection.

While previous research in SER has primarily focused on speaker personalization, this work focuses on listener personalization by adapting to previously unseen annotators. For future research, we plan to explore the applicability of Meta-PerSER to other subjective learning tasks beyond SER, such as hate speech detection and customer experience recognition. We also plan to investigate alternative meta-adaptation strategies and extend our approach to multilingual and low-resource settings [34].

7. References

- [1] B. Schuller, G. Rigoll, and M. Lang, "Speech emotion recognition combining acoustic features and linguistic information in a hybrid support vector machine-belief network architecture," in *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2004.
- [2] R. A. Khalil, E. Jones, M. I. Babar, T. Jan, M. H. Zafar, and T. Alhussain, "Speech Emotion Recognition Using Deep Learning Techniques: A Review," *IEEE Access*, 2019.
- [3] R. Nagase, T. Fukumori, and Y. Yamashita, "Speech Emotion Recognition with Fusion of Acoustic- and Linguistic-Feature-Based Decisions," in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2021.
- [4] W. Ren, Y.-C. Lin, H.-C. Chou, H. Wu, Y.-C. Wu, C.-C. Lee, H.-Y. Lee, H.-M. Wang, and Y. Tsao, "EMO-Codec: An In-Depth Look at Emotion Preservation Capacity of Legacy and Neural Codec Models with Subjective and Objective Evaluations," in *2024 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2024, pp. 1–6.
- [5] E. Mower, A. Metallinou, C.-C. Lee, A. Kazemzadeh, C. Busso, S. Lee, and S. Narayanan, "Interpreting ambiguous emotional expressions," in *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*, 2009, pp. 1–8.
- [6] V. Sethu, E. M. Provost, J. Epps, C. Busso, N. Cummins, and S. Narayanan, "The Ambiguous World of Emotion Representation," 2019. [Online]. Available: <https://arxiv.org/abs/1909.00360>
- [7] C. L. Moine, N. Obin, and A. Roebel, "Speaker Attentive Speech Emotion Recognition," in *Interspeech 2021*, 2021.
- [8] M. Tran, Y. Yin, and M. Soleymani, "Personalized Adaptation with Pre-trained Speech Encoders for Continuous Emotion Recognition," in *Interspeech 2023*, 2023.
- [9] B. Plank, "The "Problem" of Human Label Variation: On Ground Truth in Data, Modeling and Evaluation," in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 2022.
- [10] L. Aroyo, A. Taylor, M. Díaz, C. Homan, A. Parrish, G. Serapio-García, V. Prabhakaran, and D. Wang, "DICES Dataset: Diversity in Conversational AI Evaluation for Safety," in *Advances in Neural Information Processing Systems*, 2023.
- [11] V. Prabhakaran, A. Mostafazadeh Davani, and M. Diaz, "On Releasing Annotator-Level Labels and Information in Datasets," in *Proceedings of the Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, C. Bonial and N. Xue, Eds., 2021.
- [12] H. Wu, H.-C. Chou, K.-W. Chang, L. Goncalves, J. Du, J.-S. R. Jang, C.-C. Lee, and H.-Y. Lee, "Open-Emotion: A Reproducible EMO-Superb For Speech Emotion Recognition Systems," in *2024 IEEE Spoken Language Technology Workshop (SLT)*, 2024, pp. 510–517.
- [13] H.-C. Chou and C.-C. Lee, "Every Rating Matters: Joint Learning of Subjective Labels and Individual Annotators for Speech Emotion Classification," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [14] A. M. Davani, M. Díaz, and V. Prabhakaran, "Dealing with Disagreements: Looking Beyond the Majority Vote in Subjective Annotations," *Transactions of the Association for Computational Linguistics*, 2022.
- [15] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, "Meta-Learning in Neural Networks: A Survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 9, pp. 5149–5169, 2022.
- [16] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ser. ICML'17. JMLR.org, 2017, p. 1126–1135.
- [17] S. Ravi and H. Larochelle, "Optimization as a Model for Few-Shot Learning," in *International Conference on Learning Representations*, 2017.
- [18] T. Fujioka, T. Homma, and K. Nagamatsu, "Meta-Learning for Speech Emotion Recognition Considering Ambiguity of Emotion Labels," in *Interspeech 2020*, 2020, pp. 2332–2336.
- [19] S. Chopra, P. Mathur, R. Sawhney, and R. R. Shah, "Meta-Learning for Low-Resource Speech Emotion Recognition," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6259–6263.
- [20] A. M. Davani, M. Díaz, and V. Prabhakaran, "Dealing with Disagreements: Looking Beyond the Majority Vote in Subjective Annotations," *Transactions of the Association for Computational Linguistics*, vol. 10, pp. 92–110, 01 2022.
- [21] Z.-T. Liu, B.-H. Wu, M.-T. Han, W.-H. Cao, and M. Wu, "Speech emotion recognition based on meta-transfer learning with domain adaption," *Applied Soft Computing*, vol. 147, p. 110766, 2023.
- [22] R. K. Gandhi, V. Tsouvalas, and N. Meratnia, "On efficacy of Meta-Learning for Domain Generalization in Speech Emotion Recognition," in *2023 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*, 2023, pp. 421–426.
- [23] H.-C. Chou and C.-C. Lee, "Learning to Recognize Per-Rater's Emotion Perception Using Co-Rater Training Strategy with Soft and Hard Labels," in *Interspeech 2020*, 2020, pp. 4108–4112.
- [24] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "IEMOCAP: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, 2008.
- [25] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin, T.-H. Huang, W.-C. Tseng, K. tik Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. Mohamed, and H. yi Lee, "SUPERB: Speech Processing Universal Performance Benchmark," in *Interspeech 2021*, 2021, pp. 1194–1198.
- [26] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: a framework for self-supervised learning of speech representations," in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- [27] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2021.
- [28] 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.
- [29] 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)*, June 2019.
- [30] A. Antoniou, H. Edwards, and A. Storkey, "How to train your MAML," in *International Conference on Learning Representations*, 2019.
- [31] B. Shi, M. Sun, K. C. Puvvada, C.-C. Kao, S. Matsoukas, and C. Wang, "Few-Shot Acoustic Event Detection Via Meta Learning," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [32] J. Opitz and S. Burst, "Macro f1 and macro f1," *arXiv preprint arXiv:1911.03347*, 2019.
- [33] 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.
- [34] H.-C. Lin, Y.-C. Lin, H.-C. Chou, and H.-y. Lee, "Improving speech emotion recognition in under-resourced languages via speech-to-speech translation with bootstrapping data selection," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.