

# Enhancing Open-Set Speaker Identification through Rapid Tuning with Speaker Reciprocal Points and Negative Samples

*Anonymous submission to Interspeech 2024*

## Abstract

This paper introduces a novel framework for open-set speaker identification in household environments, playing a crucial role in facilitating seamless human-computer interactions. Addressing the limitations of current speaker models and classification approaches, our work integrates an AudioLM frontend with a few-shot rapid tuning neural network (NN) backend for enrollment, employing task-optimized Speaker Reciprocal Points Learning (SRPL) to enhance discrimination across multiple target speakers. Furthermore, we propose an enhanced version of SRPL (SRPL+), which incorporates negative sample learning with both speech-synthesized and real negative samples to significantly improve open-set SID accuracy. Our approach is thoroughly evaluated across various multi-language datasets, demonstrating its effectiveness in achieving high usability for complex household multi-speaker scenarios and enhancing open-set performance by up to 27% over the direct use of state-of-the-art AudioLM model.

**Index Terms:** speaker identification, speaker recognition, open-set learning, few-shot learning, speech synthesis

## 1. Introduction

In household environments, the use of AI agents to accurately identify speakers through speech, a method not limited by posture or obstacles, is increasing. This is crucial for enabling seamless interaction between humans and computers, including interactions with large language model (LLM). Speaker identification, a subtask of speaker recognition, employs learning approaches to precisely recognize in-house speakers by effectively utilizing the voice print characteristics of target speakers[1].

Recent advancements in speaker recognition have featured the development of advanced speaker models, such as TDNN-based [2], ResNet-based [3], and state-of-the-art self-supervised learning models like WavLM [4, 5] and other large audio pre-trained models (AudioLMs) [6]. Despite their general applicability in representing speakers, these models primarily focus on binary classification [7], leaving room for improvement in speaker identification (SID) algorithms for high usability in complex household multi-speaker scenarios [8].

In SID, it is assumed that test utterances originate from a set of pre-enrolled speakers, employing strategies like multi-class classification loss [9, 10], prototype learning loss for few-shot learning [11, 12], or graph-based learning method [13]. However, these approaches predominantly target *closed-set* classification, which may not adequately serve real-world applications facing the challenge of *open-set* speaker identification, where the goal is to optimize detection of target-set speakers with contrastive accuracy and robustness to outliers [3].

Recent advancements in open-set SID have explored using

multiple Probabilistic Linear Discriminant Analysis (PLDA) modules for outlier detection [14] and prototype-based loss [15, 3] for improved identification. However, these methods, not inherently designed for open-set recognition, often fail to achieve optimal performance or lack the streamlined characteristic of neural network (NN)-based approaches [15].

In response to the increasing interest in open-set learning capabilities across various fields of pattern recognition [11, 16], we propose a novel framework that integrates an AudioLM frontend with a few-shot rapid tuning NN backend for enrollment. This approach employs task-optimized Speaker Reciprocal Points Learning (SRPL) to improve open-set discrimination across multiple target speakers. Furthermore, we introduce an enhanced version of SRPL (SRPL+), incorporating negative sample learning with both speech-synthesized and real negative samples to significantly improve the accuracy of open-set SID. Our main contributions include:

- Introduction of SRPL for open-set speaker identification, facilitating rapid tuning enrollment on AudioLM.
- Incorporation of negative sample learning with SRPL+, utilizing speech-synthesized and real negative samples for enhanced identification accuracy.
- Comprehensive evaluations on various multi-language datasets, both qualitatively and quantitatively.

For detailed information on open-sourced resources and reproduction, visit our project website<sup>1</sup>.

## 2. Speaker Reciprocal Points Learning for Open-set Speaker Identification

To enhance speaker distinction in specific domains, we propose the use of lightweight neural adaptation models. These models refine speaker embeddings for improved alignment with target scenarios. Figure 1 illustrates our approach, which combines a WavLM frontend [4] and a TDNN model pretrained on speaker recognition tasks. The process generates an initial embedding  $\mathbf{Emb}_{LM}$ , which the Lightweight Adapter transforms into a domain and speaker-specific embedding  $\mathbf{Emb}_{spk}$ :

$$\mathbf{Emb}_s = \text{Adapter}(\text{TDNN}(\mathbf{Emb}_{LM})) \quad (1)$$

This efficient adaptation serves as the enrollment phase in speaker recognition.

### 2.1. Rapid AudioLM Backend Tuning Approach with Speaker Reciprocal Points Learning (SRPL)

Our approach to open-set SID leverages advanced learning strategies, emphasizing few-shot learning for rapid tuning to

<sup>1</sup><https://srplplus.github.io>

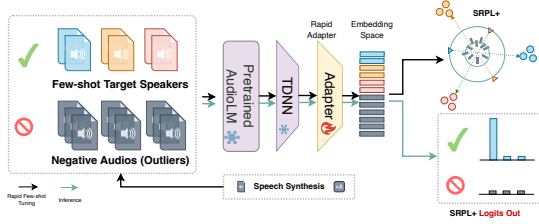


Figure 1: Illustration of open-set speaker identification architecture: customization via audio large model (LM) with SRPL-based backend rapid tuning.

enhance speaker recognition with limited data. Moving beyond traditional prototype loss methods [12, 15] and drawing inspiration from [11], we adapt reciprocal points learning for speaker recognition, termed Speaker Reciprocal Points Learning (SRPL). SRPL effectively distinguishes both known and unknown speakers, ensuring an optimal distribution of known speaker embeddings and establishing a dedicated area for characterizing unknown speakers, as shown in Figure 2. Based on generalized pretrained AudioLM models, SRPL enhances the system’s robustness for open-set recognition.

As illustrate in Figure 3, SRPL aims to maximize the distance between learnable embeddings and Reciprocal Points (RPs), with task-specific optimizations tailored for speaker recognition. Since the AudioLM utilizes angle-based optimization, diverging from the original approach [11], we calculate the distance between RPs and speaker embeddings using the *inner product*, with the goal of maximizing this distance. Consequently, the probability of identifying the target speaker is evaluated based on the distance between the adapted speaker embeddings and RPs.

$$p(k|\mathbf{Emb}_s) = \frac{e^{-\mathbf{Emb}_s \cdot \mathbf{RP}^k}}{\sum_{i=1}^K e^{-\mathbf{Emb}_s \cdot \mathbf{RP}^i}}, \quad (2)$$

$$\mathcal{L}_s(\mathbf{Emb}_{\text{LM}}, y = k; \theta) = -\log p(k|\mathbf{Emb}_{\text{LM}}), \quad (3)$$

where  $\mathbf{RP}^k$  are the RPs for the corresponding known speakers  $k$ .




Figure 2: Conceptual illustration of the embedding space for various open-set training losses. Figure is adapted from [11].

The core of SRPL focuses on optimizing the speaker manifold by confining or reserving a general unknown pool  $\mathbf{U}$  within a predefined radius  $R$ , by referencing to RPs. The method aims to ensure that the maximum distance from the unknown sample set  $\mathbf{U}$  to all reciprocal points in RP-set  $\mathcal{R}$ , does not exceed  $R$ .

$$\max(\text{Dist}(\mathbf{U}, \mathcal{R})) \leq R. \quad (4)$$

Omitting detailed deductions and proofs as outlined in [11], we can safely formulate the well-established and feasible dual training loss as a marginal-based loss.

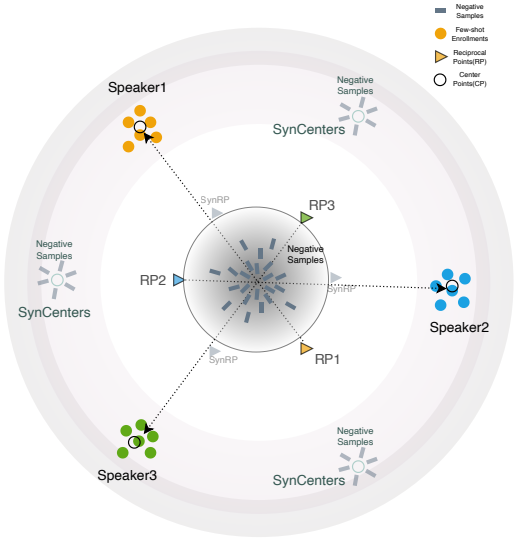


Figure 3: SRPL and its enhanced SRPL+ with integration of negative samples.

$$D_e(\mathbf{Emb}_s, \mathbf{RP}^k) = \|\mathbf{Emb}_s - \mathbf{RP}^k\|_2^2,$$

$$\mathcal{L}_r(\mathbf{Emb}_{\text{LM}}; \theta, y = k) = \max(D_e(\mathbf{Emb}_s, \mathbf{RP}^k) - R, 0), \quad (5)$$

where we utilize Euclidean distance  $D_e(\cdot, \cdot)$  to constrain the magnitude of unknowns. This optimization does not rely on prior knowledge of unknown negative data. Interestingly, the goal can be achieved by optimizing known speaker embeddings in relation to their reciprocal points.

To enhance SRPL further, we introduce center-focused learning to improve the alignment of target speaker data with its center points  $\mathbf{CP}^k$ , a well-established method for enhancing training stability [1].

$$\mathcal{L}_c(\mathbf{Emb}_{\text{LM}}, y = k; \theta) = -\log \frac{e^{\mathbf{Emb}_s \cdot \mathbf{CP}^k}}{\sum_{i=1}^K e^{\mathbf{Emb}_s \cdot \mathbf{CP}^i}}, \quad (6)$$

Putting them together, the SRPL learning algorithm is performed as:

$$\mathcal{L}_{\text{SRPL}}(\mathbf{E}_{\text{spk}}, y; \theta) = \mathcal{L}_s(\cdot) + \lambda_r \mathcal{L}_r(\cdot) + \lambda_c \mathcal{L}_c(\cdot). \quad (7)$$

## 2.2. SRPL Enhancement with Negative Samples (SRPL+)

### 2.2.1. Zero-shot Speech Synthesis Process for Negative Samples Synthesis

For typical household few-shot SID tasks, acquiring a pool of unknowns  $\mathbf{U}$  is often not feasible, leading to SRPL optimization primarily with known samples. While this approach is effective, SRPL inherently can support enhancement through the integration of negative samples that represent unknowns, further improving its performance. As suggested by [11], employing a GAN for negative sample generation is a recognized method. However, our evaluation finds such approach is resource-intensive and unstable. To streamline the process for speaker identification, we propose a zero-shot text-to-speech (ZS-TTS) synthesis module to generate multi-speaker identity speech samples, as illustrate in Figure 1.

Negative samples are generated using randomly selected reference speakers from the LibriTTS dataset [17], ensuring their voiceprints are distinct from those used in evaluations, thus forming  $\mathbf{U}_{syn}$ . Prosody controls are applied to enrich the diversity of negative examples. This approach enhances the SRPL framework, termed SRPL+. We utilize (anonymized) ZS-TTS system [X], a zero-shot TTS model trained on the LibriTTS corpus, facilitating efficient and diverse negative sample creation.

### 2.2.2. Learning Enhancement with Negative Audio Instance

For SRPL+, we refine our approach to utilize audio-level negative samples within the SRPL learning framework, as illustrated in Figure 3. This integration comprises two main processes. First, it employs negative samples to enhance performance by generating synthetic fake speakers, increasing the robustness of synthesized Reciprocal Points (SynRPs). Next, negative data primarily supports the adversarial learning process, using negative samples  $u_{syn} \in \mathbf{U}_{syn}$  to maximize their distance from all RPs, effectively targeting maximum entropy  $H(\cdot)$ . This strategy realizes the SRPL+ optimization goal, enhancing the model’s ability to distinguish between known and unknown speakers:

$$\mathcal{L}_{SRPL+}(\mathbf{Emb}_{LM}, y, \mathbf{U}_{syn}; \theta) = \mathcal{L}_{SRPL}(\cdot) - \lambda_{ns} H(u_{syn}, \mathfrak{R}). \quad (8)$$

## 3. Experiments Settings

This section details our experimental setup, datasets, metrics, and procedures, offering an overview necessary for replication. Full details are on our project website<sup>1</sup>.

### 3.1. Datasets

- **Qualcomm Speech Dataset:** This dataset contains speech utterances of four English keywords by 50 speakers, designed for evaluating our open-set SID approach in a household multi-class classification scenario [18]. We apply 3-fold cross-validation, calculating the mean performance across datasets. For each split, a selection of 10 speakers is reserved for testing targets, and 15 for testing outliers, with the remainder potentially serving as real negative samples for SRPL+. No overlap between outlier testing data and any training-stage speaker data is ensured.
- **FFSVC HiMia:** Featuring utterances of the Chinese wake-up command ‘Hi Mia’ by over 80 speakers [19], this dataset’s dev-set captures a range of distances and microphone-array recordings. We focus on the near-field portion of the FFSVC dev-set for our experiments. Similar to the Qualcomm setup, 10 speakers are reserved for target testing, employing 3-fold cross-validation. The experimental setup mirrors that used with the Qualcomm Speech Dataset.

### 3.2. Metrics

For evaluating our open-set SID tasks, we focus on multi-target classification metrics common to the patten recognition and speech processing domain. Area Under the Receiver Operating Characteristic Curve (AUROC or AUC) is used as a threshold-independent metric for SID [20]. It measures the probability that a model correctly ranks a randomly chosen target example higher than a randomly chosen negative example by plotting the true positive rate against the false positive rate.

Nonetheless, AUC may not fully capture the unique challenges of open-set recognition, which requires effectively rejecting outliers while ensuring target speakers are accurately identified. To address this, we employ the Open Set Classification Rate (OSCR) [11, 21], a more robust metric that simulates open-set recognition dynamics:

$$CCR(TH) = \frac{|\{x \in TestData^k \mid \arg \max_k P(k|x) = k \cap P(k|x) \geq TH\}|}{|TestData^k|}, \quad (9)$$

$$FPR(TH) = \frac{|\{x \mid x \in Unknown \cap \max_k P(k|x) \geq TH\}|}{|Unknown|}, \quad (10)$$

where  $TH$  is a threshold. OSCR calculates the area under the curve mapping the Correct Classification Rate (CCR) for known classes to the False Positive Rate (FPR) for unknown data, offering a threshold-independent evaluation for open-set.

### 3.3. Training and Inference Details

For our open-set SID system, we employ rapid fine-tuning with SRPL or SRPL+ in few-shot settings. During the enrollment phase, each target speaker contributing 20 enrolled utterances alongside 20 synthesized utterances for model rapid tuning. For SRPL+, we incorporate 1000+ negative samples from unknown speakers, synthesized via ZS-TTS (SynNeg) or sourced from the unused dataset portions (RealNeg). The system adapts using a 3-layer MLP, linearly transforming to K-way speaker outputs, and applies the SRPL/SRPL+ loss. Fine-tuning the adapter employs an SGD optimizer for 100 epochs, a process completed in mere seconds. Rather than using cosine similarity, we evaluate test speaker embeddings against reciprocal points using the calculated logits from Equation (2). All RPs are learnable and hyperparameter include  $\lambda_s$  are all set equally, without careful search to achieve effectiveness.

## 4. Results

### 4.1. Comparative Evaluation of SRPL with Baselines

Table 1 shows our SRPL system’s performance relative to various baseline methods. A basic comparison point is the WavLM TDNN system, which applies cosine similarity for speaker identification, using mean embeddings of enrollment audios and cosine scores for probability output across different channels. This approach reduces SID to a binary classification problem without optimizing for multi-class scenarios. It does not reasonably detects close-set scenarios, particularly with up to 10 speakers in the FFSVC evaluation, and struggles with outliers, showing its vulnerability in open-set contexts.

More advanced methodologies, such as Softmax fine-tuning [10], prototype learning [12], and OpenFEAT learning [15]—which builds on prototype learning by incorporating negative samples—are specifically designed for rapid tuning and enrollment in speaker identification. Our SRPL system surpasses these methods, including the WavLM-TDNN baseline, in both Prototype and OpenFEAT approaches. Its superiority is evident not just in close-set evaluations, where many methods excel in targeting speaker accuracy, but crucially in open-set metrics, underscoring SRPL’s robust performance in 5-way and 10-way classification tasks.

With the introduction of negative samples, SRPL+ significantly outperforms baseline methods in OSCR scores, highlighting its advancement. This enhancement is also observed with synthetic negative samples (SRPL+(SynNeg)), which

Table 1: Experimental results of open-set SID tasks.

Method	Open-set Eval		Close-set Eval	Open-set Eval		Close-set Eval
	AUC(%) $\uparrow$	OSCR(%) $\uparrow$	ACC(%) $\uparrow$	AUC(%) $\uparrow$	OSCR(%) $\uparrow$	ACC(%) $\uparrow$
Qualcomm Speech [18]	5Way			10Way		
WavLM-TDNN CosineDirect [4]	84.09	83.61	99.54	83.81	81.49	96.06
SoftmaxTune [10]	71.33	70.95	99.07	66.46	66.91	98.91
ProtoTypeTune [12]	71.54	71.39	99.54	88.87	88.45	98.91
OpenFEAT [15]	73.23	72.88	99.07	88.05	87.35	98.24
SRPL (Ours)	84.12	84.95	99.53	89.53	89.16	98.47
SRPL+(SynNeg)	92.51	92.48	<u>99.90</u>	92.06	91.25	<u>98.69</u>
SRPL+(RealNeg)	<b>95.73</b>	<b>94.25</b>	<u>99.21</u>	<b>95.40</b>	<b>94.26</b>	<u>98.81</u>
FFSVC HiMia [19]	5Way			10Way		
WavLM-TDNN CosineDirect [4]	85.07	80.27	93.50	78.61	67.24	84.25
SoftmaxTune [10]	75.87	75.86	99.80	88.66	87.60	97.25
ProtoTypeTune [12]	75.77	74.67	97.50	88.87	86.90	95.25
OpenFEAT [15]	80.10	78.98	97.00	85.63	82.83	94.75
SRPL(Ours)	89.12	88.91	99.00	91.28	90.27	98.50
SRPL+(SynNeg)	<b>95.71</b>	<b>94.97</b>	<u>99.00</u>	<b>95.48</b>	<b>94.83</b>	<u>99.00</u>
SRPL+(RealNeg)	<u>93.04</u>	<u>91.93</u>	<u>98.00</u>	<u>94.32</u>	<u>92.71</u>	<u>97.80</u>

demonstrates excellent results. Moreover, in close-set metrics, SRPL+ maintains strong performance, showcasing its comprehensive effectiveness in open-set speaker identification.

#### 4.2. Supplementary Analyses for SRPL for Open-set SID

Figure 4 showcases a t-SNE visualization of speaker embeddings, revealing that, for the WavLM-TDNN baseline, embeddings are not distinctively separated from unknowns, indicating a lack of discriminability. In contrast, SRPL’s rapid tuning markedly enhances the discrimination of speaker embeddings, effectively clustering unknowns into specific regions and improving the grouping of known target speakers. The softmax optimization, focusing on the angle of distribution, performs suboptimally compared to SRPL. While the OpenFEAT method, an advancement on prototype learning, shows improved class cohesion when learning with negatives, it falls short in separating negative instances.

Table 2 evaluates the efficacy of SRPL’s components. Utilizing an angular-aware distance metric, concentrating on center points, and incorporating negative samples for both synthesized centers and adversarial learning significantly bolster SRPL’s overall performance in speaker recognition. The backend tuning process is efficient, requiring only 200 seconds to safely converge on GPU. The use of 1000+ negative data for optimal performance adds only an additional 80 seconds, highlighting its practical feasibility.

Table 2: Ablation Studies and Training Cost Analysis for SRPL.

Method	Open-set (10-Way)		Tuning Cost
	AUC(%) $\uparrow$	OSCR(%) $\uparrow$	(GPU)
<b>SRPL+</b>	<b>95.40</b>	<b>94.26</b>	280s
w/o SynCenters	93.31	92.83	
<b>SRPL</b>	<u>89.53</u>	<u>89.16</u>	200s
w/o CenterFocus	88.05	88.30	
w/o SpkTaskOptimize	87.50	86.01	

### 5. Conclusions and Limitations

The challenges of current speaker recognition systems for open-set speaker identification are evident. This study introduces

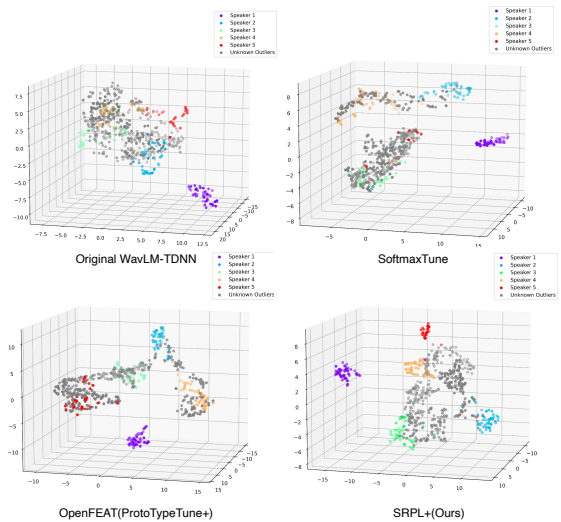


Figure 4: t-SNE visualization of speaker embeddings for targets and outliers in testing datasets for SRPL+ and baseline systems.

Speaker Reciprocal Points Learning (SRPL), a novel algorithm specifically optimized for open-set SID, leveraging the state-of-the-art AudioLM combined with a few-shot rapid tuning neural network backend for enrollment. This approach models unknown speakers using known speaker samples. It is further enhanced by SRPL+, effectively trained with negative samples, including zero-shot TTS-synthesized samples or real samples collected from household AI agents. Our method demonstrates a significant performance improvement over existing techniques, notably enhancing usability in household environments. Despite its effectiveness, we acknowledge that the open-set recognition performance is not yet optimal for some highly complex and far-field scenarios, likely due to the inherent limitations of the AudioLM backbone. Moving forward, we aim to refine our methodology to better address these challenges, potentially incorporating multimodal audio-visual approaches or other methods to enhance system robustness.

## 6. Acknowledgements

(Anonymous)

## 7. References

- [1] Z. Bai and X.-L. Zhang, "Speaker recognition based on deep learning: An overview," *Neural Networks*, vol. 140, pp. 65–99, 2021.
- [2] F. Sigona and M. Grimaldi, "Validation of an ecapa-tdnn system for forensic automatic speaker recognition under case work conditions," *Speech Communication*, p. 103045, 2024.
- [3] R. Peri, S. O. Sadjadi, and D. Garcia-Romero, "Voxwatch: An open-set speaker recognition benchmark on voxceleb," *arXiv preprint arXiv:2307.00169*, 2023.
- [4] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [5] J. Peng, O. Plchot, T. Stafylakis, L. Mošner, L. Burget, and J. Černocký, "An attention-based backend allowing efficient fine-tuning of transformer models for speaker verification," in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 555–562.
- [6] J.-w. Jung, W. Zhang, J. Shi, Z. Aldeneh, T. Higuchi, B.-J. Theobald, A. H. Abdelaziz, and S. Watanabe, "Espnet-spk: full pipeline speaker embedding toolkit with reproducible recipes, self-supervised front-ends, and off-the-shelf models," *arXiv preprint arXiv:2401.17230*, 2024.
- [7] B. Han, Z. Chen, and Y. Qian, "Exploring binary classification loss for speaker verification," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [8] A. Q. Ohi, M. F. Mridha, M. A. Hamid, and M. M. Monowar, "Deep speaker recognition: Process, progress, and challenges," *IEEE Access*, vol. 9, pp. 89 619–89 643, 2021.
- [9] R. Li, J.-Y. Jiang, C. Wu, C.-C. Hsieh, and A. Stolcke, "Speaker identification for household scenarios with self-attention and adversarial training," in *Interspeech 2020*, 2020.
- [10] Q.-B. Hong, C.-H. Wu, H.-M. Wang, and C.-L. Huang, "Combining deep embeddings of acoustic and articulatory features for speaker identification," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7589–7593.
- [11] G. Chen, P. Peng, X. Wang, and Y. Tian, "Adversarial reciprocal points learning for open set recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 11, pp. 8065–8081, 2021.
- [12] Y. Li, H. Chen, W. Cao, Q. Huang, and Q. He, "Few-shot speaker identification using lightweight prototypical network with feature grouping and interaction," *IEEE Transactions on Multimedia*, 2023.
- [13] F. Tong, S. Zheng, M. Zhang, Y. Chen, H. Suo, Q. Hong, and L. Li, "Graph convolutional network based semi-supervised learning on multi-speaker meeting data," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6622–6626.
- [14] K. Wilkinghoff, "On open-set speaker identification with i-vectors," in *Odyssey*, 2020, pp. 408–414.
- [15] K. Kishan, Z. Tan, L. Chen, M. Jin, E. Han, A. Stolcke, and C. Lee, "Openfeat: Improving speaker identification by open-set few-shot embedding adaptation with transformer," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7062–7066.
- [16] G. Chen, L. Qiao, Y. Shi, P. Peng, J. Li, T. Huang, S. Pu, and Y. Tian, "Learning open set network with discriminative reciprocal points," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 2020, pp. 507–522.
- [17] Z. Qin, W. Zhao, X. Yu, and X. Sun, "Openvoice: Versatile instant voice cloning," *arXiv preprint arXiv:2312.01479*, 2023.
- [18] Y.-H. Lee and N. Cho, "Phonmatchnet: phoneme-guided zero-shot keyword spotting for user-defined keywords," *arXiv preprint arXiv:2308.16511*, 2023.
- [19] X. Qin, M. Li, H. Bu, S. Narayanan, and H. Li, "Far-field speaker verification challenge (ffsvc) 2022: Challenge evaluation plan," 2022.
- [20] A. B. Nassif, I. Shahin, S. Hamsa, N. Nemmour, and K. Hirose, "Casa-based speaker identification using cascaded gmm-cnn classifier in noisy and emotional talking conditions," *Applied Soft Computing*, vol. 103, p. 107141, 2021.
- [21] A. R. Dhamija, M. Günther, and T. Boulton, "Reducing network agnostophobia," *Advances in Neural Information Processing Systems*, vol. 31, 2018.