

OPENFEAT: IMPROVING SPEAKER IDENTIFICATION BY OPEN-SET FEW-SHOT EMBEDDING ADAPTATION WITH TRANSFORMER

Kishan K C^{1†} Zhenning Tan^{2†} Long Chen² Minh Jin²
Eunjung Han² Andreas Stolcke² Chul Lee²

¹Rochester Institute of Technology, Rochester, NY, USA

²Amazon Alexa AI, Sunnyvale, CA, USA

ABSTRACT

Household speaker identification with few enrollment utterances is an important yet challenging problem, especially when household members share similar voice characteristics and room acoustics. A common embedding space learned from a large number of speakers is not universally applicable for the optimal identification of every speaker in a household. In this work, we first formulate household speaker identification as a few-shot open-set recognition task and then propose a novel embedding adaptation framework to adapt speaker representations from the given universal embedding space to a household-specific embedding space using a set-to-set function, yielding better household speaker identification performance. With our algorithm, **Open-set Few-shot Embedding Adaptation with Transformer** (openFEAT), we observe that the speaker identification equal error rate (IEER) on simulated households with 2 to 7 hard-to-discriminate speakers is reduced by 23% to 31% relative.

Index Terms— speaker identification, embedding adaptation, few-shot open-set learning

1. INTRODUCTION

Voice assistants, such as Alexa and Google Home, rely on speaker identification to enable personalized experiences, such as playing one’s favorite music or customizing calendars and alarms. On shared devices with multiple enrolled speakers, voice assistants need to distinguish household members based on their speech to identify the user. Since household members often share similar voice characteristics and acoustic conditions, identifying the speaker in this scenario becomes more difficult than for random groups of speakers. On real-world data, we found that the mean cosine similarity between speaker profiles and same-household utterances is about 10% greater than the similarity with utterances outside the household.

Speaker identification has two scenarios - closed-set and open-set. In closed-set speaker identification, test utterances come from one of the enrolled speakers [1]. In open-set speaker identification, the test utterance could come from a guest speaker who is not enrolled in the system [2, 3]. Open-set speaker identification is closely related to the real-world scenarios where not every utterance comes from an enrolled speaker. A typical open-set speaker identification system involves three stages. First, a universal speaker encoder is trained on a large number of speakers to generate speaker or utterance embeddings [4]. Second, a function is used to compute the distance between a test utterance embedding and each of the enrolled speaker representations. Finally, the speaker is predicted as the one that is closest to the test utterance based on the similarity measure.

The most popular speaker encoders include i-vectors [5], d-vectors [6, 7, 8], x-vectors [9, 10], and many others. The scoring function can be based on probabilistic linear discriminant analysis (PLDA) [11, 12], cosine similarity, or a neural network [13, 14].

There are two major limitations for the conventional open-set speaker recognition system. First, embeddings learned from the universal speaker encoder are not necessarily optimal for household-level speaker identification, because speech audio from household members often has shared characteristics, such as accents and room acoustics. Moreover, test utterances are compared to speaker profiles in the household *individually*, without considering the similarities shared across speakers in the household. Second, the universal speaker encoder is usually trained with classification loss or contrastive loss [4], where query utterances are mapped to one of the classes in the mini batch or episode. The model learns to divide the entire embedding space according to these classes. Unknown classes, which may lie on the decision boundaries, are not considered in this embedding space. The final speaker embeddings are forced into an embedding space that is accurate for seen speakers, which may limit the model’s power to generalize well to new speakers.

To address the first limitation, we propose an embedding adaptation by a set-to-set function. It focuses on all speaker profiles *together* in the household and transforms the universal speaker embeddings to household-specific speaker embeddings for better separation. A similar idea of embedding adaptation has been proposed in few-shot learning (FSL) tasks in the computer vision domain [15].

To address the second limitation, we train the embedding adaptation model in a FSL framework for open-set speaker identification task where guest utterances are present during episodic training. In each episode, we have query utterances from both seen classes (with support set utterances) and unseen classes (without support set utterances). For query utterances from unseen classes, we maximize the entropy of the posterior probability distribution with respect to the seen classes to enforce prediction uncertainty. We hypothesize that in this way, the model learns to map speaker embeddings conservatively, considering possible unknown speakers in the entire embedding space. This idea is first proposed in the PEELER framework in the computer vision domain [16].

To tackle both problems, we propose a new algorithm, openFEAT, to adapt speaker profile embeddings and to train the model end-to-end in a few-shot open-set recognition framework. To our knowledge, this is the first study to apply embedding adaptation to household speaker profiles to improve speaker recognition and include open-set loss to *train* speaker recognition models to improve generalization of the speaker embedding space. Our contributions are as follow: 1) We demonstrate that few-shot embedding adaptation by transformer (FEAT) and open-set recognition techniques,

[†]Equal contribution. This first author was an intern at Amazon.

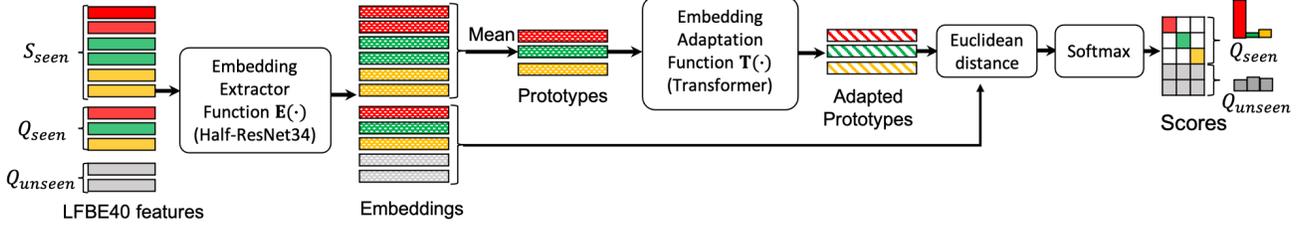


Fig. 1: Architecture of openFEAT

originally developed in the computer vision domain, can be applied to speaker recognition. 2) We evaluate both techniques and propose a combined solution (openFEAT) that can improve speaker recognition substantially.

2. PROPOSED METHOD

2.1. Few-shot open-set learning (FSOSL)

In FSL [17, 18, 19, 20], episodes (or tasks) are generated by randomly sampling a subset of the classes and a subset of examples per class. A model is trained on the episodes with classification tasks for which embeddings are optimized, producing robust embeddings (i.e., less overfitting to any particular set of classes). To train a speaker embedding extractor, in each episode D , we select N speakers with K labeled utterances per speaker, represented as an N -way K -shot problem. We denote a *support* set as $S_{\text{seen}} = \{x_i^S, y_i^S\}_{i=1}^{NK}$ where $x_i \in \mathcal{X}^S$ represents an input feature for the sample i and $y_i \in \mathcal{Y}^S$ represents its label. We denote a *seen query* set as $Q_{\text{seen}} = \{x_i^Q \in \mathcal{X}^Q, y_i^Q \in \mathcal{Y}^S\}_{i=1}^{NM}$ where M is the number of query utterances for each *seen* speaker. An embedding extractor function $E(\cdot)$ converts the input features $x_i \in \mathcal{X}^S \cup \mathcal{X}^Q$ to speaker embeddings. With the support embeddings, each class (prototype) is computed by taking an average of the class-wise embeddings: $p_c = \frac{1}{K} \sum_{x \in \mathcal{X}_c^S} E(x)$ where \mathcal{X}_c^S is a set of support instances from class c . A variety of metric-based FSL methods [21, 19, 8] use softmax classification with a distance function between the query x and the prototypes $P = (p_1, \dots, p_N)$:

$$f(x, P) = Pr(y = c|x, P) = \frac{e^{-\text{dist}(E(x), p_c)}}{\sum_{p_i \in P} e^{-\text{dist}(E(x), p_i)}} \quad (1)$$

where $Pr(\cdot)$ is a classification probability, $\text{dist}(\cdot, \cdot)$ is a temperature-scaled Euclidean distance [15], and $E(\cdot)$ is an embedding extraction function. $E(\cdot)$ is updated by cross-entropy loss as a function of class labels and the class posterior probabilities estimated by the model.

In few-shot open-set learning (FSOSL), in addition to utterances from N seen speakers ($S_{\text{seen}}, Q_{\text{seen}}$), we randomly select R *unseen* speakers, each with T utterances per episode. We denote a *unseen query* set as $Q_{\text{unseen}} = \{x_i^U \in \mathcal{X}^U, y_i^U \in \mathcal{Y}^U\}_{i=1}^{RT}$ where $\mathcal{Y}^S \cap \mathcal{Y}^U = \emptyset$. For utterances from unseen classes, the model should not assign a large classification probability to any seen classes. To enable this, a loss function should maximize the posterior entropy with respect to the seen classes [16]. This entropy loss $L_{\text{entropy}}(\cdot)$ for Q_{unseen} is combined with the cross-entropy loss $L_{\text{CE}}(\cdot, \cdot)$ for Q_{seen} to form the open-set FSL loss:

$$\mathcal{L}_{\text{open-set}} = \sum_{(x, y) \in Q_{\text{seen}}} \mathcal{L}_{\text{CE}}(y, f(x, P)) - \beta \sum_{x \in Q_{\text{unseen}}} \mathcal{L}_{\text{entropy}}(f(x, P)) \quad (2)$$

2.2. Few-shot embedding adaptation with Transformer

To adapt the prototypes to be more distinguishable in a household-specific space, we train a set-to-set function $T(\cdot)$ which transforms a set of original prototypes $P = (p_1, \dots, p_N)$ to a set of adapted prototypes $P' = (p'_1, \dots, p'_N)$ where $P' = T(P)$. We train Transformers [22, 15] as the set-to-set function:

$$\begin{aligned} \tilde{Q} &= W_Q P, \tilde{K} = W_K P, \tilde{V} = W_V P, \\ V &= \tilde{V} (\text{softmax}(\frac{\tilde{Q}^T \tilde{K}}{\sqrt{m}}))^T, \\ P' &= \text{LayerNorm}(\text{Dropout}(W_{FC} V) + P) \end{aligned} \quad (3)$$

where m is the embedding dimension, $W_Q, W_K, W_V \in R^{m \times hm}$ are learnable transform matrices, h is the number of heads, and $W_{FC} \in R^{hm \times m}$ is a fully connected, trainable weight matrix. Naturally, the self-attention mechanism without positional encoding allows Transformers to learn contextual relationships between speakers without considering their ordering, which is suitable for the speaker profile adaptation use case. This algorithm is called Few-Shot Embedding Adaptation with Transformer (FEAT) [15]. While various set-to-set functions like BiLSTM, DeepSet and Graph Convolution Networks have been studied in [15], these do not perform better than FEAT. Thus, in our paper, we only consider FEAT as the algorithm choice for our FSOSL.

With a transformation function $T(\cdot)$, the classification probability in equation (1) is now rewritten as

$$f(x, P') = Pr(y = c|x, P') = \frac{e^{-\text{dist}(E(x), p'_c)}}{\sum_{p'_i \in P'} e^{-\text{dist}(E(x), p'_i)}} \quad (4)$$

where $p'_c \in P'$ is the transformed prototype for class c . The embedding adaptation loss is applied to Q_{seen} to make sure that the seen query utterances are close to the adapted prototype of the same class and far away from those of other classes:

$$\mathcal{L}_{\text{query}} = \sum_{(x, y) \in Q_{\text{seen}}} \mathcal{L}_{\text{CE}}(y, f(x, P')) \quad (5)$$

Moreover, the embedding adaptation module should preserve the class-wise similarity after adaptation. To ensure instance embeddings after adaptation are close to their class neighbors and far away from other classes, we apply a contrastive loss as in [15]. To calculate the contrastive loss, the utterance instances from $Q_{\text{seen}} \cup S_{\text{seen}}$ are transformed by the embedding adaptation module instead of prototypes. Then, the new class centers C are calculated with the utterances in $Q_{\text{seen}} \cup S_{\text{seen}}$. The adapted instance embeddings are forced to stay close to their class center and far away from other class centers through the cross-entropy loss as:

$$\mathcal{L}_{\text{contrastive}} = \sum_{(x, y) \in Q_{\text{seen}} \cup S_{\text{seen}}} \mathcal{L}_{\text{CE}}(y, f(x, C)) \quad (6)$$

where C is a set of class centers after the transformation. Finally, the overall embedding adaptation loss is defined as:

$$\mathcal{L}_{\text{FEAT}} = \mathcal{L}_{\text{query}} + \alpha \mathcal{L}_{\text{contrastive}} \quad (7)$$

During inference, all steps are similar to those for the typical speaker identification system, except that the adapted speaker profiles used for comparison in each household are pre-computed by the embedding adaptation function $T(\cdot)$. Therefore, the embedding adaptation module does not add any run-time computation to the standard speaker identification system.

2.3. openFEAT

With the previously defined embedding adaptation, speaker profiles are now mapped to a new space. To ensure the newly adapted speaker profiles do not collide with unknown guest speaker embeddings in the new space, we propose to add query utterances from Q_{unseen} during embedding adaptation in FSOSL in a similar fashion as that of [16]. We call our FSOSL algorithm as openFEAT (Figure 1). Our openFEAT is based on the intuition that the newly adapted embeddings are only adjusted locally to improve the separability between speaker sets and should not move to a completely different embedding space so as to cause collisions with unseen speakers. The new total loss for openFEAT is:

$$\mathcal{L}_{\text{openFEAT}} = \mathcal{L}_{\text{query}} + \alpha \mathcal{L}_{\text{contrastive}} - \beta \sum_{x \in Q_{\text{unseen}}} \mathcal{L}_{\text{entropy}}(f(x, P')) \quad (8)$$

where α and β are hyper-parameters that control the weight of different losses. Note that the entropy loss for Q_{unseen} is computed with the adapted prototypes P' .

3. EXPERIMENTS

We used VoxCeleb datasets in our experiments. VoxCeleb2 dev [7] is used to train the embedding function $E(\cdot)$ and the embedding adaptation module $T(\cdot)$. The VoxCeleb1 dataset [23] is used for evaluation. There is no overlap between speakers in these two datasets.

3.1. Household simulation for evaluation

To evaluate household level speaker identification, we simulated households with different sizes (n). Speakers in the same households may share similar characteristics. To mimic this real-life difficulty, we synthesize households with hard-to-discriminate speakers based on the cosine similarity between the speaker embeddings. Utterance level embeddings are generated using the embedding function $E(\cdot)$ from pretrained models as in Section 3.2, and speaker level embeddings are constructed by averaging up to 100 randomly selected utterances. We then select highly similar speakers, using the 85th percentile among all cosine similarity scores between speaker profiles as a threshold.

During evaluation, for each speaker, 4 utterances are used as enrollment utterances to create average speaker enrollment profiles, and 10 random utterances are selected as evaluation utterances. In addition, we randomly select $50 \times n$ (where n is the household size) utterances from other speakers outside the household as guest utterances for evaluation. We want the number of guest utterances to be large enough to give reliable performance estimates. With the number chosen here, results for simulated households have small 95% confidence intervals, as shown in Table 1. For each enrollment, evaluation or guest utterance, we sample 10 equally spaced 3-second segments and average their embeddings to represent the utterance.

For each household size, we run the simulation five times and report the 95% confidence interval together with the mean of the evaluation metric.

3.2. Pretraining the embedding extractor

For better initialization for embedding adaptation in FSL, we first pretrain the embedding extractor $E(\cdot)$ on VoxCeleb2 dev data with 800,000 episodes. We chose a half-ResNet34 [24], with half the channels in the residual block as the original ResNet34 [25], and with a self-attentive pooling (SAP) layer as $E(\cdot)$. The output embedding dimension is 256. In each episode, we selected $N = 400$ speakers, each with $M = 1$ support utterance and $K = 1$ query utterance, as this was the largest episode that could fit into the GPU memory. We adopted prototypical loss [19] to learn the embedding extraction function $E(\cdot)$. An Adam optimizer with initial learning rate 10^{-3} was used and the learning rate was decayed by a factor of 0.95 every 10 epochs. We used $\frac{1}{32}$ as temperature to scale the similarity measure $\text{dist}(\cdot, \cdot)$. This pretrained embedding extractor $E(\cdot)$ was then used to extract embeddings in our household simulation. Our embedding extractor has 3.2% equal error rate (EER) on the VoxCeleb1 test set, comparable to what was reported in [4].

3.3. Fine-tune embedding extractor

In baseline experiments (column ‘‘Baseline’’ in Table 1), where only unadapted embeddings are used, we fine-tuned the embedding extractor on the same dataset as used in pretraining using episodes that are close to practical household sizes, for 16000 episodes. For each episode, we selected $N = 10$, which is close to, but larger than, typical household sizes so as to ensure sufficient training task difficulty. We set $K = 4$ to mimic four enrollment utterances per speaker. Each speaker had $Q = 5$ as query utterances. For open-set FSL, we selected an additional $R = 5$ speakers, each with $T = 5$ query utterances without support utterances. Optimizer, learning rate and its decay schedule are the same as in pretraining.

3.4. Train embedding adaptation

The embedding adaptation module $T(\cdot)$ was jointly trained with the embedding extractor $E(\cdot)$ using the same episodes as for fine-tuning the embedding extractor. $E(\cdot)$ was initialized with the pretrained weights, while $T(\cdot)$ was randomly initialized. As the embedding adaptation module we use a Transformer as in [15], using a self-attention mechanism with single head and single layer. A dropout rate of 0.5 is applied after a fully connected layer, as in [15]. The set-to-set function $T(\cdot)$ does not change the embedding dimension. Therefore, the output dimension of $T(\cdot)$ is 256 as well.

3.5. Model evaluation

During evaluation, we scored utterance similarity using cosines scaled to the interval $[0, 1]$. For evaluation utterances in each household, we computed the false acceptance rate (FAR) and false negative identification rate (FNIR). FAR is the fraction of guest utterances that are falsely accepted as enrolled speakers. FNIR is the fraction of enrolled speaker utterances that are misidentified as a different speaker or rejected as a guest speaker. We aggregate false acceptance errors and false negative identification errors and define the identification equal error rate (IEER) as the point when FAR equals FNIR. Thus, our error metric is similar to the EER used in speaker verification, but takes the error types of the speaker identification task into account.

4. RESULTS

4.1. Embedding adaptation and FSOSL

Table 1: IEER (%) of embedding adaptation models on simulated households (lower is better). n is household size. Numbers in parentheses show relative IEER improvement over baseline. We used $\mathcal{L}_{\text{FEAT}}$ with $\alpha = 0.5$ for FEAT, $\mathcal{L}_{\text{open-set}}$ with $\beta = 0.1$ for open-set FSL, and $\mathcal{L}_{\text{openFEAT}}$ with $\alpha = 0.5, \beta = 0.1$ for openFEAT.

n	Baseline	FEAT	Open-set	openFEAT
2	6.48±0.29	4.91±0.15(24.3%)	5.16±0.31(20.4%)	4.49±0.20(30.7%)
3	8.65±0.14	6.75±0.12(22.0%)	7.06±0.21(18.4%)	6.06±0.18(30.0%)
4	10.56±0.26	8.56±0.12(18.9%)	8.73±0.15(17.4%)	7.67±0.21(27.4%)
5	11.98±0.18	10.01±0.23(16.5%)	10.04±0.26(16.2%)	9.02±0.24(24.8%)
6	13.46±0.12	11.37±0.12(15.5%)	11.23±0.18(16.5%)	10.30±0.23(23.5%)
7	14.69±0.37	12.45±0.38(15.3%)	12.35±0.35(16.0%)	11.35±0.37(22.8%)

We evaluated the performance on simulated households and observed that baseline IEER increases with household size, which is expected given that speaker identification becomes harder the more speakers are enrolled. When evaluating the FEAT model, we observe relative improvements of 15.3% to 24.3% across all households. For FSOSL, the IEER improvement is 16.0% to 20.4% across all households. When combining embedding adaptation with FSOSL, the IEER is reduced further by 22.8% to 30.7% relative.

It is noteworthy that the relative improvement decreases with increasing household size. Since embedding adaptation works by projecting the embeddings to a new space while keeping clusters more separable, we hypothesize that with more speakers in the household, finding such a perfect projection space becomes more difficult; therefore, the improvement levels off with increasing household size.

We also conducted ablation studies to optimize the hyperparameters α and β that control the strength of contrastive learning loss and open-set loss. We find that $\alpha = 0.5$ achieves the best embedding adaptation performance for FEAT. Considering both hyperparameters, openFEAT with $\alpha = 0.5$ and $\beta = 0.1$ achieves the best performance.

4.2. Embedding visualization

To qualitatively evaluate the effect of the FEAT embedding adaptation module, we visualized speaker profiles before and after embedding adaptation. We first learn a household-specific PCA projection module using the preadaptation embeddings of the support and query sets for each household. Then, we apply the learned PCA projection to speaker profiles before and after adaptation, along with the evaluation utterance embeddings, both from within and outside the household. As shown in Figure 2, the adapted speaker profiles are further apart from each other. Interestingly, with openFEAT, we find the speaker profiles can be better separated from guest utterances even along the first two dimensions of the PCA projection.

4.3. Score distribution

We also evaluated changes in the utterance score distribution after speaker profile adaptation. For test utterances from enrolled speaker A, we compute the cosine score between each utterance and its corresponding speaker profile A (AA), as well as to the closest other speaker profile B (AB). For guest utterances (G), randomly selected from outside the household, we only compute the cosine score to the closest speaker profile B (GB) in the household. As shown in Figure 3, with embedding adaptation (FEAT), the score distributions have all shifted to a lower score range and the overlap between them is reduced, reflecting lower probability of error. With openFEAT, the

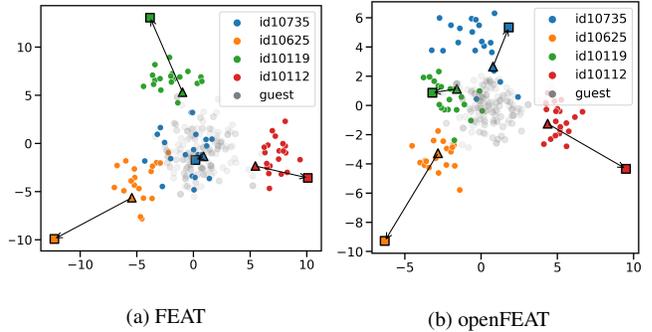


Fig. 2: PCA visualization of test utterance embeddings and speaker profiles before and after adaptation in a four-speaker household. Dots represent utterance embeddings. Colored dots are from enrolled speakers, while grey dots are from guest utterances randomly selected from outside the household. Triangles and squares represent the speaker profiles before and after adaptation.

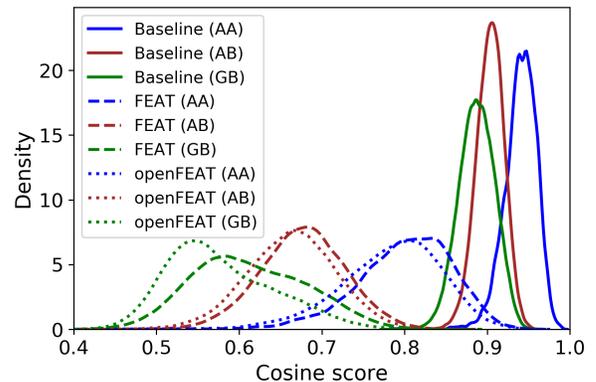


Fig. 3: Test utterance score distributions for four-speaker households. AA represents the score between an utterance from enrolled speakers and its true speaker. AB represents the score between an utterance from enrolled speakers and its closest other speaker in the household. GB represents the score between a guest utterance and its closest speaker in the household.

guest utterance score distribution is further shifted to a lower range, indicating that the model learns an embedding space for better separation between guest utterances and household speakers.

5. CONCLUSIONS

We developed an embedding adaptation method for speaker identification in a household setting, based on Transformer-based set-to-set mapping of profile embeddings. We further incorporated guest utterances into training episodes to reduce speaker embedding confusability with unseen speakers in FSOSL. Our proposed method openFEAT achieves relative IEER reduction of 23% to 31% for simulated households of hard-to-discriminate speakers on VoxCeleb1 dataset. Furthermore, we observe that embedding adaptation trained with open-set loss achieves better separation of speaker profiles both in PCA visualizations and score distributions.

6. ACKNOWLEDGMENTS

We thank our Speaker Understanding team and Aparna Khare for their valuable feedback, and managers for their inputs and support.

7. REFERENCES

- [1] R. Li, J.-Y. Jiang, X. Wu, C.-C. Hsieh, and A. Stolcke, "Speaker Identification for Household Scenarios with Self-Attention and Adversarial Training," in *Proc. Interspeech*, 2020, pp. 2272–2276.
- [2] J. Fortuna, P. Sivakumaran, A. Ariyaceinia, and A. Malegaonkar, "Open-set speaker identification using adapted Gaussian mixture models," in *Proc. 9th European Conference on Speech Communication and Technology*, 2005, pp. 1997–2000.
- [3] K. Wilkinghoff, "On open-set speaker identification with ivectors," in *Proc. Odyssey Speaker and Language Recognition Workshop*, 2020, pp. 408–414.
- [4] J. S. Chung, J. Huh, S. Mun, M. Lee, H. S. Heo, S. Choe, C. Ham, S. Jung, B.-J. Lee, and I. Han, "In defence of metric learning for speaker recognition," in *Proc. Interspeech*, 2020, pp. 2977–2981.
- [5] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.
- [6] G. Heigold, I. Moreno, S. Bengio, and N. Shazeer, "End-to-end text-dependent speaker verification," in *Proc. IEEE ICASSP*, 2016, pp. 5115–5119.
- [7] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Interspeech*, 2018, pp. 1086–1090.
- [8] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *Proc. IEEE ICASSP*, 2018, pp. 4879–4883.
- [9] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification," in *Proc. Interspeech*, 2017, pp. 999–1003.
- [10] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. IEEE ICASSP*, 2018, pp. 5329–5333.
- [11] S. Ioffe, "Probabilistic linear discriminant analysis," in *Proc. European Conference on Computer Vision*. Springer, 2006, pp. 531–542.
- [12] S. J. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *Proc. IEEE 11th International Conference on Computer Vision*, 2007, pp. 1–8.
- [13] D. Garcia-Romero, G. Sell, and A. McCree, "Magneto: X-vector magnitude estimation network plus offset for improved speaker recognition," in *Proc. Odyssey Speaker and Language Recognition Workshop*, 2020, pp. 1–8.
- [14] J. Pelecanos, Q. Wang, and I. L. Moreno, "Dr-vectors: Decision residual networks and an improved loss for speaker recognition," in *Proc. Interspeech*, 2021, pp. 4603–4607.
- [15] H.-J. Ye, H. Hu, D.-C. Zhan, and F. Sha, "Few-shot learning via embedding adaptation with set-to-set functions," in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8808–8817.
- [16] B. Liu, H. Kang, H. Li, G. Hua, and N. Vasconcelos, "Few-shot open-set recognition using meta-learning," in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8798–8807.
- [17] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," in *Advances in Neural Information Processing Systems*, vol. 29, 2016, pp. 3630–3638.
- [18] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," in *Proc. International Conference on Learning Representations*, 2017.
- [19] J. Snell, K. Swersky, and R. S. Zemel, "Prototypical networks for few-shot learning," *Advances in Neural Information Processing Systems*, vol. 30, pp. 4077–4087, 2017.
- [20] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. International Conference on Machine Learning*. PMLR, 2017, pp. 1126–1135.
- [21] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 815–823.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [23] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Proc. Interspeech*, 2017, pp. 2616–2620.
- [24] H. S. Heo, B.-J. Lee, J. Huh, and J. S. Chung, "Clova baseline system for the VoxCeleb speaker recognition challenge 2020," *arXiv preprint arXiv:2009.14153*, Sep. 2020.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.