



# Meta-Learning for Short Utterance Speaker Recognition with Imbalance Length Pairs

Seong Min Kye<sup>1</sup>, Youngmoon Jung<sup>1</sup>, Hae Beom Lee<sup>1</sup>, Sung Ju Hwang<sup>1,2</sup>, Hoirin Kim<sup>1</sup>

<sup>1</sup>KAIST, South Korea

<sup>2</sup>AITRICS, South Korea

{kye9165, dudans, haebeom.lee, sjhwang82, hoirkim}@kaist.ac.kr

## Abstract

In practical settings, a speaker recognition system needs to identify a speaker given a short utterance, while the enrollment utterance may be relatively long. However, existing speaker recognition models perform poorly with such short utterances. To solve this problem, we introduce a meta-learning framework for imbalance length pairs. Specifically, we use a Prototypical Networks and train it with a support set of long utterances and a query set of short utterances of varying lengths. Further, since optimizing only for the classes in the given episode may be insufficient for learning discriminative embeddings for unseen classes, we additionally enforce the model to classify both the support and the query set against the entire set of classes in the training set. By combining these two learning schemes, our model outperforms existing state-of-the-art speaker verification models learned with a standard supervised learning framework on short utterance (1-2 seconds) on the VoxCeleb datasets. We also validate our proposed model for unseen speaker identification, on which it also achieves significant performance gains over the existing approaches. The codes are available at <https://github.com/seongmin-kye/meta-SR>.

**Index Terms:** speaker verification, speaker identification, meta-learning, short duration, text-independent, open-set

## 1. Introduction

Speaker recognition (SR) with short utterance is an important challenge since in practical settings, where test utterance could be as short as 1 to 5 seconds. In the past decades, i-vector [1] combined with probabilistic linear discriminant analysis (PLDA) [2, 3] has been a dominant approach for SR. However, recently, deep neural network (DNN) based methods have shown to outperform i-vector based systems, achieving better performance on short utterances [4] over non-deep learning counterparts. Although the recent advances in deep learning make it possible to obtain impressive performance on SR such as speaker verification (SV) and identification (SI), obtaining sufficiently high performances under practical settings (e.g. short duration SR, unseen speaker SI) remains as a challenging problem. Some recent works propose to tackle these issues. Regarding SR with short utterances, [5–7] introduce task-specific feature extractor which allows to extract more information from short utterances. [6–8] introduce aggregation methods to attend to more informative frames in frame-level features. In addition to these approaches, various attempts have been made to deal with short utterances. However, they do not provide a substantial performance improvements under realistic scenarios.

In this work, we aim to tackle this problem by meta-learning with imbalance length pairs. Specifically, we organize each episode such that it contains a support set of long utterances and a query set of variable short utterances. By optimizing sequence of episodes, we can train our network to match long-short utterance

pairs well over conventional (referred to as ‘vanilla’) training, which optimizes for the same-length utterances. Also, variable length of query utterances allows the model to consider various practical situations during meta-training and thus allows us to obtain a more length-robust model for SR.

Yet, a crucial problem here is that the query samples may not be discriminative against the entire set of classes (speakers) in the training set. Thus, we further classify every sample in each episode against the whole training classes (referred to as ‘global classification’). In doing so, an embedding of a short utterance becomes discriminative against other classes and is matched to its own long utterance at the same time (See Figure 1). Also, having a consistent framework across training and test phases, by targeting unseen speakers during training, allows the model to verify and identify unseen speakers well (See Figure 2).

Our proposed learning scheme uses a ResNet34 [9] as the base network architecture, which is widely used for SR. To verify the efficacy of our proposed model, we use a simplistic framework for implementation of our model, using naive pooling methods such as Temporal Average Pooling (TAP) for aggregation and non-margin metric loss. Also, we use 40-dimensional log mel-filterbank features as inputs to reduce time complexity, since the execution time should be short in realistic settings. We experiment on various settings, such as short utterance SV and unseen SI including conventional experimental settings (full utterance SV). We use VoxCeleb datasets [10, 11] to directly compare with other models. Our model obtains the state-of-the-art results on short utterance SV and unseen SI on various datasets, although we use simpler implementations with smaller-dimensional features.

Our main contributions are as follows:

- We propose a meta-learning framework for short utterance speaker recognition, in which each episode is composed of support and query pairs with imbalance length utterances.
- We further propose a training objective that combines the episodic classification loss with the global classification loss, which allows to obtain well-matched and discriminative embeddings.
- We validate our model on VoxCeleb datasets under various realistic scenarios, including speaker verification with short duration and unseen speaker identification, and achieve the state-of-the-art results.

## 2. Related Work

**DNN based speaker embedding:** Recently, DNN based methods [12–17], have achieved impressive performance on speaker recognition (SR), outperforming traditional i-vector systems. The key components of DNN based systems are feature extractors, aggregation of temporal features and optimization. First,




Figure 1: Comparison between (a) vanilla training and (b) meta-learning with global classification. We visualize  $t$ -SNE embeddings on 10-way 5-shot task. Prototype is average of support set which consists of 5 seconds utterances. Best viewed in color.

many SR systems use 1D or 2D convolutional neural networks and recurrent neural networks as feature extractors, which make it possible to extract the time and frequency properties of the speaker features (MFCC, mel-filterbank). Then these extracted frame-level features are summarized into fixed-length vectors using aggregation methods which aims to capture intrinsic speaker information, such as attentive statistic pooling (ASP) [18], self attentive pooling (SAP) [19], learnable dictionary encoding (LDE) [19] and spatial pyramid encoding (SPE) [20]. Then, these utterance-level features are used as inputs for with softmax classifiers with fully-connected layers. However, since softmax classifier may not obtain sufficiently discriminative embedding spaces, recent methods such as A-softmax [21], AM-softmax [22] and AMM-softmax [23] propose angular margin-based metric to reduce per-class variance.

**Metric-based meta-learning for few-shot classification:** Speaker verification and speaker identification for unseen speaker are essentially few-shot learning tasks. The goal of few-shot classification is to correctly classify unlabeled query (test) examples with only a few labeled support (enrollment) set per class. Since labeled data is scarce, conventional supervised learning in this case is prone to overfitting, and thus recent works resort to *meta-learning*, which learns a model that generalize over diverse tasks to consider larger number of problems for training. One of the most popular meta-learning approach is metric-based meta-learning [24–27], which learns an embedding space that is learned to be discriminative for any given tasks. Several recent works [28, 29] use Prototypical Networks [26], which is such a metric-based meta-learning model, for speaker recognition. Compared with above methods, our algorithm uses not only Prototypical Networks with imbalance length pairs but also global classification over the whole samples in episode.

**Speaker recognition for short utterance:** Speaker recognition for short utterance is especially challenging since the input data contains very little information about the speaker. To tackle this problem, [5–8] propose aggregation techniques to extract as much information as possible from short speech. NetVLAD / GhostVLAD [8] use attention-based pooling with learnable dictionary encoding, and time-distributed voting (TDV) [6] utilizes short-cut connection information with weighted sum of them, which yields impressive performance gains over GhostVLAD for short utterance SR. However, TDV obtains good performance only on short utterances and relatively low performance on long utterances. Other approaches to tackle the short utterance SR use knowledge distillation [30], generative adversarial networks [31] and angular margin-based method [32, 33].

### 3. Method

In this work, we consider a practical setting for unseen speaker recognition, where the length of test utterance is shorter than the

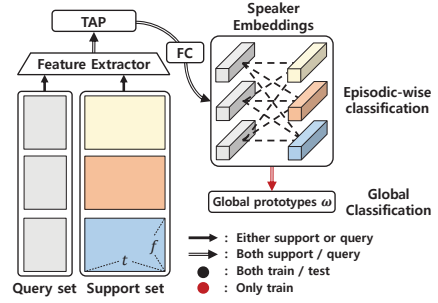


Figure 2: Overview of proposed meta-learning scheme on 3-way 1-shot task. We denote the features for each speaker using different colors.

enrollment utterance. To solve this problem, the model should not only match a pair of utterances with different lengths from the same speaker, but also be able to recognize unseen speakers not included in training set well. To this end, we introduce a metric-based meta-learning scheme, in which support and query sets consist of long and short utterances, respectively. We also classify both support and query sets over the entire set of training classes rather than training them only for the given set of classes to obtain even more discriminative embeddings.

#### 3.1. Problem Definition

The goal of few-shot unseen short utterance speaker recognition problem is to recognize test utterance  $\tilde{x}_i$  which is as short as 5 seconds as  $i$ -th speaker, given only few utterances  $x_i$  from each speaker for enrollment, whose lengths could be longer than 5 seconds. Since the number of enrollment examples per class is too small, conventional supervised learning may obtain suboptimal performance due to overfitting.

Thus, we tackle the problem with episodic meta-learning where we learn a model over diverse tasks, such that the model learns to recognize the speaker for *any* utterances, while considering a different classification problem at each time. To this end, we first compose task episodes with support set and query set. We first randomly sample  $N$  classes from the given dataset, and then sample  $K$  and  $M$  examples from each class as the support set and query set, respectively. We define the task sampling distribution as  $p(\tau)$ . As a result, we have a support set  $\mathcal{S} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N \times K}$  and a query set  $\mathcal{Q} = \{(\tilde{\mathbf{x}}_i, \tilde{y}_i)\}_{i=1}^{N \times M}$ , where  $y, \tilde{y} \in \{1, \dots, N\}$  are the class labels.

#### 3.2. Meta-learning with imbalance length pairs

Despite large progress on speaker recognition, speaker recognition with short utterances remains to be very challenging in realistic settings due to length mismatch between training and test utterances. As shown in [34], in conventional training setting, this can result in performance degeneration for short utterances. A prior work [33] tackles this problem by training the model with short segments only, which outperforms models with long utterances, but the model obtains relatively poor performances on long speech. This trade-off is problematic in realistic settings where the test utterances could be given in any lengths.

How can we then train a length-robust model for speaker recognition? To tackle this problem, we construct the training episodes for our meta-learning framework, to contain imbalanced length pairs. In practical settings, while we can enroll long utterances to the system, the test utterances could come in any variable lengths. To simulate this situation in training phase, we set the length of the utterance for the support set to be longer

than the query utterances. On the other hand, we construct the query sets such that they have variable lengths (1 to 2 seconds), that are smaller than the lengths of the support utterances.

As with [26], we compute the class prototypes by averaging over the support set and enforce the query examples to become closer to their own prototypes. First, we define  $\mathcal{S}_c$  as the set of support examples in class  $c$  and then compute the prototype of each class  $c = 1, \dots, N$  in episode:

$$P_c = \frac{1}{|\mathcal{S}_c|} \sum_{x \in \mathcal{S}_c} f_\theta(\mathbf{x}) \quad (1)$$

Then, we compute the distance between each query and its corresponding prototype. In this work, we use cosine similarity as the distance metric:

$$d(f_\theta(\tilde{\mathbf{x}}_i), P_c) = \frac{f_\theta(\tilde{\mathbf{x}}) \cdot P_c}{\|P_c\|_2} = \|f_\theta(\tilde{\mathbf{x}})\|_2 \cdot \cos(\theta_{i,c}) \quad (2)$$

where it can be seen as the cosine similarity with an input-wise length scale. We can then obtain the probability of the sample belonging to each class  $c$  as follows:

$$p(\tilde{y} = c | \tilde{\mathbf{x}}, \mathcal{S}; \theta) = \frac{\exp(d(f_\theta(\tilde{\mathbf{x}}), P_c))}{\sum_{c'=1}^C \exp(d(f_\theta(\tilde{\mathbf{x}}), P_{c'}))} \quad (3)$$

Then, we compute the loss for each episode:

$$L_e^\tau(\theta) = \frac{1}{|Q|} \sum_{(\tilde{\mathbf{x}}, \tilde{y}) \in Q} -\log p(\tilde{y} | \tilde{\mathbf{x}}, \mathcal{S}; \theta) \quad (4)$$

### 3.3. Global classification

With the proposed meta-learning scheme, we can make variable short utterance to be close to relatively long utterance. However, optimization only within each episode may fail to learn a discriminative embedding space. Inspired by [24], we additionally classify support and query samples against whole training classes. By optimizing support and query samples of different length at once, we can reduce per class variance according to utterance duration. At the same time, we can make the discriminative embeddings over all other classes in training set. Following [24], we assume a set of global prototypes for each class:

$$\omega = \{\mathbf{w}_c \in \mathbb{R}^l | c = 1, \dots, C'\} \quad (5)$$

where  $C'$  is the number of classes in entire training set and  $l$  is dimension of embedding. For  $(\mathbf{x}, \mathbf{y}) \in S \cup Q$ , we predict the probability of the sample  $x$  being an instance of class  $y$  as follows:

$$p(\mathbf{y} | \mathbf{x}; \theta, \omega) = \frac{\exp(d(f_\theta(\mathbf{x}), \mathbf{w}_y))}{\sum_{c=1}^{C'} \exp(d(f_\theta(\mathbf{x}), \mathbf{w}_c))} \quad (6)$$

and compute the global loss:

$$L_g^\tau(\theta, \omega) = \frac{1}{|S| + |Q|} \sum_{(\mathbf{x}, \mathbf{y}) \in S \cup Q} -\log p(\mathbf{y} | \mathbf{x}; \theta, \omega) \quad (7)$$

where  $d$  is the distance metric described in Eq. 2. Note that the global classification is conducted on both the support and query samples. Finally, our learning objective combines the episode loss in Eq. 4 with the global loss in Eq. 7.

$$L(\theta, \omega) = \mathbb{E}_{p(\tau)} [L_e^\tau(\theta) + \lambda L_g^\tau(\theta, \omega)] \quad (8)$$

Here,  $\lambda$  is the hyperparameter for loss balancing and we simply use  $\lambda = 1$ . In order to compute the final objective, we sample a single task and then average loss according to the task distribution  $p(\tau)$  during training. This combined objective allows our model to match imbalance length pairs, while these pairs are classified over whole training classes together.

Table 1: Verification performance on full utterance. *C*: Contrastive loss; *A*: A-softmax; *SM*: Softmax; *NS*: Normalized softmax; *M*: Meta-learning; *Spec*: Spectrogram; \*: With data augmentation; *D*: Development set.

Model (Aggregation+Loss)	M	Feature	Train dataset	$C_{det}^{min}$ / EER%
i-vectors+PLDA [10]	✗	–	Vox1(D)	0.73 / 8.8
VGG-M (TAP+C) [10]	✗	Spec-512	Vox1(D)	0.71 / 7.8
ResNet34 (SAP+A) [19]	✗	MFB-64	Vox1(D)	0.622 / 4.40
ResNet34 (SPE+A) [20]	✗	MFB-64	Vox1(D)	<b>0.402</b> / 4.03
TDNN (ASP+SM) [18]	✗	MFCC-40	Vox1(D)*	0.406 / 3.85
<b>ResNet34 (TAP+NS)</b>	✓	MFB-40	Vox1(D)	0.418 / <b>3.81</b>
UtterIdNet (TDV+SM) [6]	✗	Spec-257	Vox2(D)	– / 4.26
Thin ResNet34 [8] (GhostVlad+SM)	✗	Spec-257	Vox2(D)	– / 3.22
ResNet34 (SPE+A) [20]	✗	MFB-64	Vox2(D)	0.245 / 2.61
<b>ResNet34 (TAP+NS)</b>	✓	MFB-40	Vox2(D)	<b>0.234</b> / <b>2.08</b>

## 4. Experiments

### 4.1. Dataset

We experiment our method under various settings on VoxCeleb datasets. VoxCeleb1 [10] and VoxCeleb2 [11] are large scale text-independent speaker recognition datasets, each of which consists of 1251 and 5994 speakers, respectively. The two datasets have disjoint sets of speakers. We measure the speaker verification results with equal error rate (EER) and the minimum detection cost function (minDCF or  $C_{det}^{min}$ ) at  $P_{target} = 0.01$ . We score the verification trials using cosine similarity. For unseen speaker identification, we report the average accuracy over 1000 randomly generated episodes with 95% confidence intervals.

### 4.2. Experiment setting

We use 40-dimensional log mel-filterbank (MFB) features with a frame-length of 25 ms as input features, which overlaps the adjacent frames by 15ms. We mean-normalized the inputs along the time-axis without any voice activity detection (VAD) or data augmentation. In training episodes, we perform 1-shot 100-way classification, while setting the number of query examples for each class to 2. For memory efficiency, we set length of support set to 2 seconds and the length of the query to between half and full of the support length. For vanilla training, we use fixed length speech of 2 seconds. For frame-level feature extraction, we use ResNet34 with 32-64-128-256 channels for each residual stage. Extracted features are aggregated with TAP and are passed through the fully-connected layer to obtain 256-dimensional embeddings. We use SGD optimizer with the Nesterov momentum of 0.9 and set the weight decay to 0.0001. We set initial learning rate to 0.1 and decay it by a factor of 10 until convergence. Every experiment is done with a single NVIDIA 2080Ti GPU.

### 4.3. Speaker verification for full utterance

We first examine the results of full-duration SV to analyze the advantage of our training scheme. The results in Table 1 show the model performances evaluated on VoxCeleb1 [10] original test trial. For fair comparison, we report baselines without VAD and data augmentation except for x-vector [14] based models [18]. On VoxCeleb1, our proposed model outperforms previous state-of-the-arts models. For the same backbone (i.e. ResNet34), our model achieves superior performance without any aggregation and margin-based metrics. In general, additional aggregation and margin-based metric lead to the better performance. Further, our model outperforms time delay neural network (TDNN) with attentive statistic pooling [18]. Our model also consistently outperforms baseline models, on a larger VoxCeleb2 [11] dataset.

Table 2: Verification performance on short utterance. G: Global classification; M: Meta-learning; D: Development set; T: Test set.

Model (Aggregation+Loss)	G	M	Feature	Train dataset	Test dataset	EER%	EER%	EER%
						1s	2s	5s
ResNet34 (TAP+NS)	✓	✗	MFB-40	Vox1(D)	Vox1(T)	9.41	6.76	5.22
ResNet34 (TDV+NS) [6]	✓	✗	MFB-40	Vox1(D)	Vox1(T)	9.12	6.44	4.92
ResNet34 (TAP+NS)	✗	✓	MFB-40	Vox1(D)	Vox1(T)	8.44	6.33	4.81
<b>ResNet34 (TAP+NS)</b>	✓	✓	MFB-40	Vox1(D)	Vox1(T)	<b>7.53</b>	<b>5.39</b>	<b>4.03</b>
Thin ResNet34 (GhostVlad+SM) [8]	✓	✗	Spec-257	Vox2(D)	Vox1(D+T)	12.71	6.59	3.34
ResNet34 (SAP+AM) [33]	✓	✗	MFB-80	Vox2(D)*	Vox1(D+T)	9.91	4.48	2.26
<b>ResNet34 (TAP+NS)</b>	✓	✓	MFB-40	Vox2(D)	Vox1(D+T)	<b>5.31</b>	<b>3.15</b>	<b>2.17</b>

Table 3: Performance comparison on length pairs of training utterances. All models are meta-learned without global classification. VoxCeleb1 development set and test set are used for training and testing, respectively.

Support length	Query length	EER%	EER%	EER%	EER%
		1s	2s	5s	full
2s	2s	9.32	6.66	5.05	4.64
2s	1s	8.81	6.56	5.36	4.94
2s	1s-2s	<b>8.44</b>	<b>6.33</b>	<b>4.81</b>	<b>4.52</b>

#### 4.4. Speaker verification for short utterance

We first describe the experimental settings, then report the results of our model and other previous state-of-the-arts models for short utterance. We test our model on two datasets: 1) The original VoxCeleb1 test trial which is the same dataset used to evaluate full utterances, and 2) VoxCeleb1 full dataset (1251 speakers in total). We use full-duration enrollment utterances, but randomly cropped the test utterances by 1, 2 and 5 seconds. If a test utterance is shorter than required we set it to the target length by duplicating its own segment.

To show the efficacy of our method, we perform an ablation study with VoxCeleb1 on upper rows of Table 2. We observe that TDV [6] which aims at short segments outperforms temporal average pooling with slight margin. However, the result in the third row shows that model only trained with meta-learning outperforms TDV and conventionally trained model (See first row). Further, our proposed model which combines meta-learning with global classification obtains the best performance against other baselines trained on VoxCeleb1 with large margin.

For the comparison against other previous state-of-the-arts models [6, 8, 33], we trained the model with VoxCeleb2 dataset and tested on VoxCeleb1 full dataset. We use the same trial as described in [33]. For every speaker, we randomly generate trials for 100 positive pairs and 100 negative pairs. Bottom rows in Table 2 show that our model outperforms baseline models with significantly large margin for 1-2 seconds. Since our model does not use any aggregation techniques or margin-based optimization, we can say that its impressive improvement mostly comes from our combined learning scheme. Furthermore, note that our model uses only 40-dimensional features but the baselines use features with more than twice the dimensions. Thus our model may obtain larger performance if we use higher dimensional inputs. For the comparison against [6], since UtterIdNet is not publicly available, we compared it using TDV instead.

Our performance gain is due to two reasons. First, we compose training episode with imbalance length pairs, where utterance length of query largely varies and is shorter than the length of the utterances in the support set. In Table 3, we observe that the proposed imbalance length pairs setting outperforms both equal length pairs and fixed long-short pairs. Note that [28, 29] use equal-length pairs. In our proposed setting, model comes across various length pairs at each episode, and then is meta-learned such that it can well-match the imbalanced length pairs and become robust to speech duration. Secondly, to learn more

Table 4: Accuracy (%) of unseen speaker identification.

Query length	Training method	5-way	20-way	50-way	100-way
1s	Vanilla	94.77 $\pm$ 0.33	85.63 $\pm$ 0.29	77.72 $\pm$ 0.22	70.76 $\pm$ 0.17
	<b>Ours</b>	<b>96.40<math>\pm</math>0.30</b>	<b>88.92<math>\pm</math>0.25</b>	<b>82.10<math>\pm</math>0.20</b>	<b>75.92<math>\pm</math>0.16</b>
2s	Vanilla	97.18 $\pm$ 0.27	92.18 $\pm$ 0.23	86.63 $\pm$ 0.19	81.46 $\pm$ 0.15
	<b>Ours</b>	<b>98.38<math>\pm</math>0.20</b>	<b>94.91<math>\pm</math>0.19</b>	<b>90.90<math>\pm</math>0.16</b>	<b>86.90<math>\pm</math>0.13</b>
5s	Vanilla	98.27 $\pm$ 0.24	94.95 $\pm$ 0.18	91.30 $\pm$ 0.16	87.61 $\pm$ 0.13
	<b>Ours</b>	<b>99.12<math>\pm</math>0.16</b>	<b>97.15<math>\pm</math>0.14</b>	<b>94.88<math>\pm</math>0.12</b>	<b>92.34<math>\pm</math>0.10</b>

discriminative embeddings, we classify both the support and the query samples against the entire set of training classes. Unlike the conventional method which classifies the utterance of same length for each batch, our combined scheme classifies utterances of different lengths at once. It results in reduction of variance caused by speech duration and enhances inter-class clustering. By combining these two components, our proposed model achieves the state-of-the-arts performance on short utterances, while yielding good performance on full utterances.

#### 4.5. Unseen speaker identification

We now evaluate the performance of our model on unseen speaker identification tasks. To analyze our model, we trained the model on VoxCeleb2 dataset and tested it on the whole VoxCeleb1 dataset. As done in the verification experiments, we enroll with one utterance for each speaker and set the enrollment utterances equally to 5 seconds. Therefore, we randomly sample  $N$ -speakers from VoxCeleb1 dataset, and then sample 1 and 5 utterances from each speaker for enrollment and test utterance, respectively. For utterance shorter than required, we handled it as done in 4.4. As shown in Table 4, our proposed method outperforms vanilla training in every setting. The performance gap increases as the number of speakers grows. Generally, the performance of identification decreases as the number of speakers becomes larger and the length of utterance becomes shorter.

## 5. Conclusion

We proposed a novel meta-learning scheme for short duration speaker recognition. In order to simulate practical settings in training, we propose an episode composition in which the support and query set have different speech lengths, and combined the meta-learning scheme with global classification for obtaining discriminative embedding space. We validate our model on various speaker recognition tasks on VoxCeleb datasets, and obtain the state-of-the-art performances on short utterance speaker recognition.

## 6. Acknowledgements

This material is based upon work supported by the Ministry of Trade, Industry and Energy (MOTIE, Korea) under Industrial Technology Innovation Program (No.10063424, Development of distant speech recognition and multi-task dialog processing technologies for in-door conversational robots).

## 7. References

- [1] 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.
- [2] D. Garcia-Romero and C. Y. Espy-Wilson, "Analysis of i-vector length normalization in speaker recognition systems," in *Proc. of Interspeech*, 2011, pp. 249–252.
- [3] S. J. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *Proc. of ICCV*, 2007, pp. 1–8.
- [4] G. Bhattacharya, M. J. Alam, and P. Kenny, "Deep speaker embeddings for short-duration speaker verification," in *Proc. of Interspeech*, 2017.
- [5] Z. Gao, Y. Song, I. V. McLoughlin, W. Guo, and L.-R. Dai, "An improved deep embedding learning method for short duration speaker verification," in *Proc. of Interspeech*, 2018, pp. 3578–3582.
- [6] A. Hajavi and A. Etemad, "A deep neural network for short-segment speaker recognition," in *Proc. of Interspeech*, 2019, pp. 2878–2882.
- [7] Y. Jung, S. M. Kye, Y. Choi, M. Jung, and H. Kim, "Improving multi-scale aggregation using feature pyramid module for robust speaker verification of variable-duration utterances," *arXiv preprint arXiv:2004.03194*, 2020.
- [8] W. Xie, A. Nagrani, J. S. Chung, and A. Zisserman, "Utterance-level aggregation for speaker recognition in the wild," in *Proc. of ICASSP*, 2019, pp. 5791–5795.
- [9] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of CVPR*, 2016, pp. 770–778.
- [10] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," in *Proc. of Interspeech*, 2017, pp. 2616–2620.
- [11] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *Proc. of Interspeech*, 2018, pp. 1086–1090.
- [12] E. Varnani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification," in *Proc. of ICASSP*, 2014, pp. 4052–4056.
- [13] C. Li, X. Ma, B. Jiang, X. Li, X. Zhang, X. Liu, Y. Cao, A. Kannan, and Z. Zhu, "Deep speaker: an end-to-end neural speaker embedding system," *arXiv preprint arXiv:1705.02304*, 2017.
- [14] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *Proc. of ICASSP*, 2018, pp. 5329–5333.
- [15] C. Zhang, K. Koishida, and J. H. Hansen, "Text-independent speaker verification based on triplet convolutional neural network embeddings," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 9, pp. 1633–1644, 2018.
- [16] Y. Kwon, S.-W. Chung, and H.-G. Kang, "Intra-class variation reduction of speaker representation in disentanglement framework," *arXiv preprint arXiv:2008.01348*, 2020.
- [17] H. Han, S.-W. Chung, and H.-G. Kang, "Mirnet: Learning multiple identities representations in overlapped speech," *arXiv preprint arXiv:2008.01698*, 2020.
- [18] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive statistics pooling for deep speaker embedding," in *Proc. of Interspeech*, 2018, pp. 2252–2256.
- [19] W. Cai, J. Chen, and M. Li, "Exploring the encoding layer and loss function in end-to-end speaker and language recognition system," in *Proc. of Speaker Odyssey*, 2018.
- [20] Y. Jung, Y. Kim, H. Lim, Y. Choi, and H. Kim, "Spatial pyramid encoding with convex length normalization for text-independent speaker verification," in *Proc. of Interspeech*, 2019, pp. 4030–4034.
- [21] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, "Sphereface: Deep hypersphere embedding for face recognition," in *Proc. of CVPR*, 2017, pp. 212–220.
- [22] F. Wang, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face verification," *Proc. of IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.
- [23] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proc. of CVPR*, 2019, pp. 4690–4699.
- [24] S. M. Kye, H. B. Lee, H. Kim, and S. J. Hwang, "Meta-learned confidence for few-shot learning," *arXiv preprint arXiv:2002.12017*, 2020.
- [25] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," in *Proc. of NeurIPS*, 2016, pp. 3630–3638.
- [26] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proc. of NeurIPS*, 2017, pp. 4077–4087.
- [27] Y. Liu, J. Lee, M. Park, S. Kim, E. Yang, S. Hwang, and Y. Yang, "Learning to propagate labels: Transductive propagation network for few-shot learning," in *Proc. of ICLR*, 2019.
- [28] J. Wang, K.-C. Wang, M. T. Law, F. Rudzicz, and M. Brudno, "Centroid-based deep metric learning for speaker recognition," in *Proc. of ICASSP*, 2019, pp. 3652–3656.
- [29] P. Anand, A. K. Singh, S. Srivastava, and B. Lall, "Few shot speaker recognition using deep neural networks," *arXiv preprint arXiv:1904.08775*, 2019.
- [30] J.-w. Jung, H.-S. Heo, H.-j. Shim, and H.-J. Yu, "Short utterance compensation in speaker verification via cosine-based teacher-student learning of speaker embeddings," in *Proc. of ASRU*, 2019, pp. 335–341.
- [31] J. Zhang, N. Inoue, and K. Shinoda, "I-vector transformation using conditional generative adversarial networks for short utterance speaker verification," in *Proc. of Interspeech*, 2018, pp. 3613–3617.
- [32] Z. Huang, S. Wang, and K. Yu, "Angular softmax for short-duration text-independent speaker verification," in *Proc. of Interspeech*, 2018, pp. 3623–3627.
- [33] A. Gusev, V. Volokhov, T. Andzhukhaev, S. Novoselov, G. Lavrentyeva, M. Volkova, A. Gazizullina, A. Shulipa, A. Gorlanov, A. Avdeeva *et al.*, "Deep speaker embeddings for far-field speaker recognition on short utterances," in *Proc. of Speaker Odyssey*, 2020.
- [34] A. Kanagasundaram, S. Sridharan, S. Ganapathy, P. Singh, and C. B. Fookes, "A study of x-vector based speaker recognition on short utterances," in *Proc. of Interspeech*, 2019, pp. 2943–2947.