



# Improved Relation Networks for End-to-End Speaker Verification and Identification

Ashutosh Chaubey, Sparsh Sinha, Susmita Ghose

Data Science, LG Ads Solutions, Mountain View, California, USA

ashutoshchaubey@lgads.tv, sparshsinha@alphonso.tv, susmita@alphonso.tv

## Abstract

Speaker identification systems in a real-world scenario are tasked to identify a speaker amongst a set of enrolled speakers given just a few samples for each enrolled speaker. This paper demonstrates the effectiveness of meta-learning and relation networks for this use case. We propose improved relation networks for speaker verification and few-shot (unseen) speaker identification. The use of relation networks facilitates joint training of the frontend speaker encoder and the backend model. Inspired by the use of prototypical networks in speaker verification and to increase the discriminability of the speaker embeddings, we train the model to classify samples in the current episode amongst all speakers present in the training set. Furthermore, we propose a new training regime for faster model convergence by extracting more information from a given meta-learning episode with negligible extra computation. We evaluate the proposed techniques on VoxCeleb, SITW and VCTK datasets on the tasks of speaker verification and unseen speaker identification. The proposed approach outperforms the existing approaches consistently on both tasks.

**Index Terms:** speaker verification, speaker identification, meta-learning, relation networks

## 1. Introduction

Speaker identification and verification systems are widely used in practical settings for authentication, security, and personalized recommendations. The goal of speaker identification is to identify the speaker present in a given test sample amongst the enrolled speakers, given a few voice samples for each speaker in the enrollment set. Speaker verification models, on the other hand, try to classify whether speakers in a given pair of test clips are the same. For speaker verification, traditionally, a frontend model such as i-vector [1, 2] is used to extract speaker embeddings. A backend model is then used to compare embeddings from different voice samples [3, 4]. Embeddings corresponding to each enrolled speaker are compared with the test sample embeddings for speaker identification. The speaker with maximum similarity is identified as the speaker in the test sample.

In recent years, deep learning based speaker verification systems have gained traction, achieving superior performance compared to the traditional methods [5]. Most of these networks are trained on speaker classification tasks with additional objectives to reduce intra-class variance and increase inter-class separability [6]. For final classification, a naive cosine similarity metric is typically used and works well in most cases. Several backend classification models [7, 8] have also been proposed for this purpose, but most of these models are not trained end-to-end with the frontend.

More recently, meta-learning based approaches have proven to be very efficacious for learning good speaker embeddings for identification and verification [9]. Prototypical

network loss for speaker verification outperformed the state-of-the-art triplet loss [10, 11]. Prototypical networks have also been used for short utterance speaker verification, with an additional global classification loss for classifying samples present within an episode amongst all speakers present in the training data [12]. Additional contrastive loss and transformation coefficients were introduced to improve the quality of learned speaker embeddings [13]. These techniques use meta-learning for learning only a discriminative frontend for extracting speaker embeddings.

Inspired by the effectiveness of relation networks in few-shot image classification tasks [14], we propose relation networks for speaker identification and verification. Instead of using cosine similarity to compare two speaker embeddings in a meta-learning episode, we propose using a *relation network* that computes the similarity between two speaker embeddings. Additionally, we improve upon the input to the relation network inspired by the b-vector approach [7]. Inspired by [12], we include extra supervision to the model in terms of global classification to learn discriminative embeddings.

Meta-learning-based speaker recognition techniques converge slower than vanilla speaker classification-based techniques [15]. To improve that, we introduce a novel yet simple training regime that utilizes information from the samples present in a single meta-learning episode more efficiently.

In summary, major contributions of this work are (i) improved relation network-based model for speaker verification and identification which can be trained end-to-end; (ii) introduction of global classification with relation networks, in addition to local classification, for better supervision during training; and (iii) a new training regime which leads to faster convergence. To the best of our knowledge, this work is the first detailed analysis of the effectiveness of relation networks for speaker verification. Furthermore, we report the performance of our approach on VoxCeleb [16, 17], SITW [18] and VCTK [19] datasets and demonstrate better performance.

The rest of this paper is organized as follows. Section 2 introduces meta-learning approaches for speaker verification and relation networks. Section 3 describes the proposed approach in detail. In Section 4, we detail the experimental setup and report the results. Finally, we provide a conclusion to this work in Section 5.

## 2. Related works

### 2.1. Meta-learning

Meta-learning is one of the most common approaches for few-shot learning tasks [20]. In an  $N$ -way  $k$ -shot classification scenario, training occurs episodically. Each mini-batch consists of a *support* (enrollment) set  $S$  with  $k$  labelled samples for each of the  $N$  classes, and an unlabelled *query* (test) set  $Q$ .  $(x_{c,i}^s, y_{c,i}^s)$  represents the  $i^{th}$  support sample for class  $c$  and

Table 1: Summary of the speaker encoder network.  $T$ : temporal dimension of feature maps,  $s_k$ : kernel size,  $s_t$ : stride,  $d$ : dilation,  $n_f$ : number of features.

Layer/Block	Composition	Input size	Output size
conv1	1D Convolution $s_k = 5; s_t = 2$	80 x T	1024 x T
conv2	Dilated SE Res2Block $s_k = 3; d = 2$	1024 x T	1024 x T
conv3	Dilated SE Res2Block $s_k = 3; d = 3$	1024 x T	1024 x T
conv4	Dilated SE Res2Block $s_k = 3; d = 4$	1024 x T	1024 x T
conv5	1D Convolution $s_k = 1; s_t = 2$	3072 x T	1536 x T
atten	Attentive Stat. Pool. [15]	1536 x T	3072 x 1
FC	FC; $n_f = 192$	3072 x 1	192 x 1

$(x_j^q, y_j^q)$  represents the  $j^{\text{th}}$  query sample (here  $i \in \{1, \dots, k\}$ ;  $j \in \{1, \dots, |Q|\}$ ;  $c, y_{c,i}^s, y_j^q \in \{1, \dots, N\}$ ).

Prototypical networks [10, 21] for speaker verification use a speaker encoder (frontend)  $f_\theta$  to encode each of the samples present in the current episode to an  $M$ -dimensional speaker embedding. For each of the  $N$  speaker classes present in the current episode, prototypes  $\nu_c$  are computed by averaging the speaker embeddings of the  $k$  support samples for class  $c$ . The probability distribution  $P$  over speaker classes for a given query sample  $x^q$  is written as a softmax over the similarities  $\mathcal{S}$  between the speaker embedding  $f_\theta(x^q)$  of the test sample with all the prototypes of  $N$  speaker classes.

$$P(y^q = c|x^q) = \frac{\exp(\mathcal{S}(f_\theta(x^q), \nu_c))}{\sum_{c'=1}^N \exp(\mathcal{S}(f_\theta(x^q), \nu_{c'}))} \quad (1)$$

where  $y^q \in \{1, \dots, N\}$  is the label for sample  $x^q$ .

Relation networks [14] for few-shot learning use a relation network  $g_\phi$  (backend) for computing the *relation*  $r$  between two embeddings. Given embeddings  $f_\theta(x_i)$  and  $f_\theta(x_j)$  corresponding to a pair of samples, the relation network provides a similarity score or *relation*  $r_{i,j} \in [0, 1]$  for that pair given by,

$$r_{i,j} = g_\phi(\mathcal{C}[f_\theta(x_i), f_\theta(x_j)]) \quad (2)$$

where  $\mathcal{C}[\cdot, \cdot]$  represents a concatenation of the embeddings channel-wise. A relation network thus provides a scope for learning the similarity between a pair of embeddings.

## 2.2. DNN-based speaker encoder

Inspired from [15], we use a TDNN-based speaker encoder with channel- and context-dependent statistics pooling and 1D Squeeze Excitation (SE) Res2Blocks. Table 1 summarizes the architecture of the speaker encoder. The first layer is a 1D convolutional layer with a kernel size of 5 and a stride of 2 with 1024 channels.

The first layer is followed by three subsequent dilated SE Res2Blocks [22]. Each dilated SE Res2Block consists of a Res2Block [22] preceded and followed by a 1D convolutional layer with kernel size one. Finally, there is a squeeze excitation (SE) block which introduces channel interdependencies within each dilated SE Res2Block [23]. All the layers involved in the dilated SE Res2Blocks have 1024 channels. The multi-level

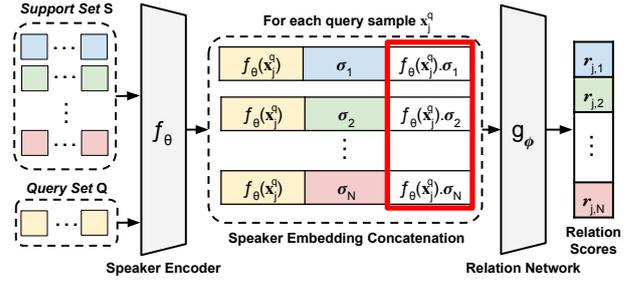


Figure 1: Schematic diagram of the proposed relation network based approach. Red box indicates the additional element-wise multiplication term input to the relation network. Colours in the support set represent different classe  $c \in \{1, \dots, N\}$ .

features obtained from the three dilated SE Res2Blocks are then aggregated through a 1D convolutional layer.

The frame-level embeddings are aggregated using channel- and context-dependent attentive statistics pooling [15] to include attention to both channel and temporal dimensions of the feature map. The aggregated feature is then passed through a fully connected layer to get the utterance-level speaker embedding.

## 3. Method

### 3.1. Improved relation networks

In this work, given a query sample  $x_j^q$  in an episode, we propose to compute the relation score between  $x_j^q$  and support samples  $x_{c,i}^s$  for each of the  $c \in \{1, \dots, N\}$  classes as depicted in Figure 1. First, we pass each sample present in an episode through the speaker encoder  $f_\theta$  to get the corresponding speaker embeddings. For each class, we average out the embeddings of the support samples to generate an aggregated representation  $\sigma_c$  of that speaker.

$$\sigma_c = \frac{1}{|S_c|} \sum_{i=1}^{|S_c|} f_\theta(x_{c,i}^s) \quad (3)$$

For computing the relation  $r_{j,c}$  between embedding pair  $f_\theta(x_j^q)$  and  $\sigma_c$ , instead of just concatenating  $f_\theta(x_j^q)$  and  $\sigma_c$  as in Eqn. (2), we propose an additional concatenation of element-wise multiplication of the two speaker embeddings (red box in Figure 1), i.e.

$$r_{j,c} = g_\phi(\mathcal{C}[f_\theta(x_j^q), \sigma_c, f_\theta(x_j^q) \cdot \sigma_c]) \quad (4)$$

where  $\mathcal{C}[\cdot, \cdot, \cdot]$  represents a concatenation of the embeddings channel-wise. Using this additional multiplication operation helps in identifying whether the embeddings correspond to the same speaker or not [6, 7].

As suggested in [14], we use mean square error (MSE) as the local loss  $\mathcal{L}_{local}$  for training our model, regressing the score  $r_{j,c}$  to the ground truth, i.e.

$$\mathcal{L}_{local} = \sum_{j=1}^{|Q|} \sum_{c=1}^N (r_{j,c} - \mathbf{1}(y_j^q == c))^2 \quad (5)$$

where  $y_j^q \in \{1, \dots, N\}$  represents the ground truth for the sample  $x_j^q$ , and  $\mathbf{1}(\cdot)$  is 1 if the condition inside bracket is true, else 0.

### 3.2. Global classification

Inspired from [12], we provide extra supervision by using global classification. We assume  $\mathcal{W} = \{\omega_C | C \in \{1, \dots, N'\}\}$  as representations or *prototypes* for all the  $N'$  classes present in the entire dataset. Then the relation  $r_{j,C}$  between the query sample  $x_j^q$  and representation  $\omega_C$  is,

$$r_{j,C} = g_\phi(C[f_\theta(x_j^q), \omega_C, f_\theta(x_j^q) \cdot \omega_C]) \quad (6)$$

The corresponding global loss function  $\mathcal{L}_{global}$

$$\begin{aligned} \mathcal{L}_{global} = & \sum_{j=1}^{|Q|} \sum_{C=1}^{N'} (r_{j,C} - \mathbf{1}(Y_j^q == C)) \\ & + \sum_{i=1}^{|S|} \sum_{C=1}^{N'} (r_{i,C} - \mathbf{1}(Y_i^s == C)) \quad (7) \end{aligned}$$

where  $Y_i^s, Y_j^q \in \{1, \dots, N'\}$  are the global labels for the sample  $x_i^s$  and  $x_j^q$ . Note that here, we are computing the loss for both the support and query samples.

The relation loss function  $\mathcal{L}_{total}$  including the global loss is given by,

$$\mathcal{L}_{total} = \mathcal{L}_{local} + \lambda * \mathcal{L}_{global} \quad (8)$$

where  $\lambda$  is a hyperparameter.

Training the relation network happens in stages. First, we train just using  $\mathcal{L}_{local}$  (Eqn. (5)), optimizing the relation network through episodic supervision. Then,  $\omega_C$  for all the classes in the training set is initialized as the average of speaker embeddings  $f_\theta(x_{C,i})$  for all the samples belonging to that class in the training set. The network is then fine-tuned with added global supervision with the objective mentioned in Eqn. (8).

### 3.3. Improved training regime

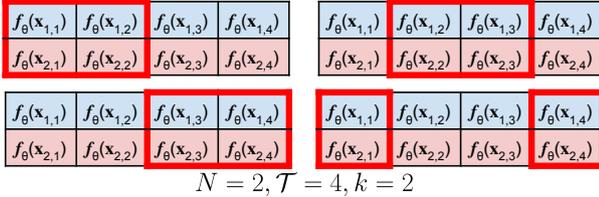


Figure 2: *Different support-query combinations. The red box represents the support set, and the rest of the embeddings are in the query set. Colours represent different classes. Instead of using the first  $k$  embeddings from each class as the support set, we use  $k$  consecutive embeddings in cyclic order.*

For a class  $c$  present in the current episode, the support samples  $x_{c,i}^s$  and query samples  $x_{c,j}^q$  are chosen randomly and can be used interchangeably. Since obtaining the speaker embeddings through the speaker encoder  $f_\theta$  is the bottleneck for computation in the training pipeline, we propose using samples present in the current episode by creating multiple support-query combinations as depicted in Figure 2.

For a given episode, we create  $\mathcal{T}$  support-query combinations, where  $\mathcal{T}$  is the total number of samples for each class within that episode. Given speaker embeddings  $F_c = \{f_\theta(x_{c,1}), \dots, f_\theta(x_{c,\mathcal{T}})\}$  for each of the  $N$  classes in the current episode, we propose creating multiple support-query com-

binations  $X_{s,q}$  as follows,

$$X_{s,q} = \bigcup_{l=1}^{\mathcal{T}} \{(S_l = \{S_{l,1}, \dots, S_{l,N}\}, Q_l = \{Q_{l,1}, \dots, Q_{l,N}\})\} \quad (9)$$

$$S_{l,c} = \{\mathcal{H}(c, l), \dots, \mathcal{H}(c, l + (k - 1))\} \quad (10)$$

$$Q_{l,c} = \{\mathcal{H}(c, l + k), \dots, \mathcal{H}(c, l + (\mathcal{T} - 1))\} \quad (11)$$

$$\mathcal{H}(c, z) = f_\theta(x_{c,1+(z-1)\%T}) \quad (12)$$

where  $(S_l, Q_l)$  is the  $l^{th}$  support-query combination,  $(S_{l,c}, Q_{l,c})$  are the support and query samples for a class  $c$  in the  $l^{th}$  support-query combination,  $c \in \{1, \dots, N\}$ .  $\%$  denotes a modulo operation and  $\mathcal{H}(c, z)$  is the cyclic index function with the property  $\mathcal{H}(c, 1) = f_\theta(x_{c,1}), \dots, \mathcal{H}(c, \mathcal{T}) = f_\theta(x_{c,\mathcal{T}}), \mathcal{H}(c, \mathcal{T} + 1) = f_\theta(x_{c,1}), \dots, \mathcal{H}(c, 2\mathcal{T}) = f_\theta(x_{c,\mathcal{T}})$  and so on.

Losses for each of the support-query combinations within an episode are aggregated and backpropagated together, only once. Note that there are more such support-query combinations possible but we only use  $\mathcal{T}$  such pairs obtained from Eqn (9) as above for computational simplicity.

## 4. Experiments

### 4.1. Datasets

We use the VoxCeleb1 [16], SITW [18] and VCTK [19] datasets to evaluate the proposed methods on the tasks of speaker verification and few-shot (unseen) speaker identification. We use the dev set of VoxCeleb2 [17] for training all the model variations. Since neural networks benefit from data augmentation, four different types of augmentations are added to the training clips following the Kaldi recipe [24] in combination with the MUSAN (music, babble, noise) [25] and RIR (reverb) [26] dataset.

For speaker verification, we use the standard trial pairs for the VoxCeleb1 test set [16] and eval *core-core* trial pairs for the SITW dataset [18] for evaluating our model. We use Equal Error Rate (EER) and the minimum detection cost function (minDCF or  $C_{det}^{min}$ ) with  $P_{target} = 0.01$  and  $C_{FA} = C_{Miss} = 1$  for comparing the proposed methods with the baselines. For unseen speaker identification, we evaluate the model using the VCTK corpus [19]. We run 1000 randomly generated episodes and report the average accuracy with 95% confidence intervals.

### 4.2. Experimental setup

We use 80-dimensional log mel-filterbanks with a window size of 25ms and a hop of 10ms as input features. Input features are mean-normalized across the time axis, and no voice activity detection (VAD) is applied. As an additional augmentation step, SpecAugment [27] is applied by randomly masking 0-10 frames in the time axis and 0-8 frames in the frequency axis.

All variations in the models along with the baseline models have been trained with an initial learning rate of 0.001 with a step-wise decay of 0.97 every epoch. We use Adam optimizer with  $\beta = \{0.9, 0.999\}$  and a weight-decay of  $2e-5$ . The meta-learning-based approaches have been trained by keeping 120 speakers in each episode with one support sample and two query samples for each speaker. For a fair comparison, we train the vanilla methods with a batch size of 360. All the relation networks are fully connected networks. We use leaky ReLU activations for better gradient flow and dropout in between the layers of the relation network to avoid overfitting. A single NVIDIA GeForce RTX 3090 GPU was used to run all our experiments.

Table 2: Performance of various systems for speaker verification. *S*: softmax, *AAM*: AAM-softmax [28], *M*: meta-learning.

System Configuration			VoxCeleb1		SITW	
Model + Loss	Backend	M	EER%	$C_{det}^{min}$	EER%	$C_{det}^{min}$
TDNN + S	cosine	✗	2.466	0.2623	4.733	0.6858
TDNN + AAM	cosine	✗	1.262	<b>0.1517</b>	2.735	0.3633
TDNN + AAM	Relation Net	✗	1.452	0.2484	3.235	0.4139
TDNN + Prototypical [12]	cosine	✓	1.235	0.1961	2.277	0.3934
TDNN + Relation Vanilla ( <b>Ours</b> )	Relation Net	✓	1.225	0.1548	2.237	0.3328
TDNN + Relation Improved ( <b>Ours</b> )	Relation Net	✓	<b>1.193</b>	0.1523	<b>2.050</b>	<b>0.3187</b>

Table 3: Performance of the proposed approach with added global supervision.  $\lambda$ : weight hyperparameter for global classification in Eqn. (8).

System Configuration		VoxCeleb1		SITW	
Model	$\lambda$	EER%	$C_{det}^{min}$	EER%	$C_{det}^{min}$
TDNN + Relation	0	1.543	0.2157	2.570	0.4744
TDNN + Relation	0.5	1.410	0.1681	2.214	0.4168
TDNN + Relation	1.0	1.336	0.1688	2.324	0.4103

### 4.3. Speaker verification

Table 2 summarizes the results of the proposed approach on speaker verification. We report the performances for two variations of our relation network-based model. The first variation uses relation as described in Eqn. (2) while the other uses the additional multiplicative term described in Eqn. (4). We can see that the improved relation network performs better than the other methods on both VoxCeleb1 and SITW. Meta-learning based approaches in general show superior performance as compared to the speaker classification-based approaches. A backend with architecture same as relation network (Eqn. (4)) was trained with the embeddings of TDNN + AAM frontend separately. Better performance over this setting indicates that the proposed end-to-end meta-learning is better than training the frontend and backend models separately. Moreover, the superior performance of the relation network-based approaches over prototypical networks highlight the importance of increased flexibility in terms of having a learnable backend.

### 4.4. Global classification

To evaluate and prove the efficacy of the added global supervision introduced in Section 3.2, we trained the proposed approach by varying  $\lambda$  in Eqn (8) and report the performance of the trained models in Table 3. All the model variations are trained on VoxCeleb2. Each episode has 30 speakers having one support and one query sample. As we increase the value of  $\lambda$ , we see a significant increase in the model’s performance. This gain in performance can be attributed to the additional supervision that is provided by global classification. These results validate that global classification improves the overall discriminability of the model.

### 4.5. Improved training regime

To analyze our training technique, we trained the relation network model on VoxCeleb2 from the same starting point using three different approaches. The vanilla approach trains the relation network-based model using the embeddings of samples in

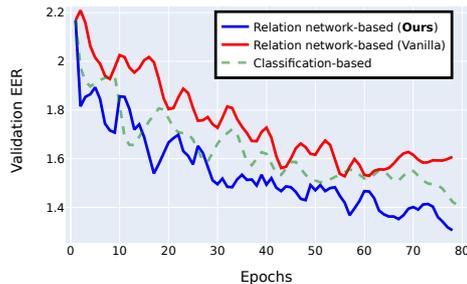


Figure 3: Training curves for different training techniques. Avg. training time per epoch - Relation network-based(**Ours**): 25.05 min, Relation network-based(Vanilla): 23.97 min, Classification based: 24.94 min.

Table 4: Unseen speaker identification performance on VCTK dataset. *AAM*: AAM-softmax [28], *M*: meta-learning.

System Configuration			Accuracy (%)		
Model + Loss	Backend	M	10-way	50-way	100-way
TDNN + AAM	cosine	✗	98.84 $\pm$ 0.07	96.10 $\pm$ 0.05	93.98 $\pm$ 0.03
TDNN + Prototypical[12]	cosine	✓	98.34 $\pm$ 0.10	95.18 $\pm$ 0.08	92.75 $\pm$ 0.07
TDNN + Relation Improved ( <b>Ours</b> )	Relation Net	✓	<b>98.98<math>\pm</math>0.05</b>	<b>96.94<math>\pm</math>0.04</b>	<b>95.24<math>\pm</math>0.04</b>

the current episode just once, while our approach uses the embeddings more efficiently (refer Section 3.3). We also trained a classification-based TDNN + AAM model to compare the meta-learning-based approaches with it. Figure 3 shows the training curve for the different approaches for 80 epochs. Validation EERs are reported on trial pairs constructed from a small subset (2%) of VoxCeleb2 disjoint from the training set. We can observe that our improved training regime leads to faster convergence. Both the vanilla and our improved technique have almost equal computation times as the time to process the embeddings through the relation network is negligible compared to generating the embeddings from the speaker encoder.

### 4.6. Unseen speaker identification

Unseen speaker identification means that none of the speakers used in testing has been used while training the model. We evaluate the performance of unseen speaker identification on VCTK [19] corpus with one support and five query samples in each test episode. As shown in Table 4, the proposed relation network-based approach performs better than the TDNN + Prototypical [12] as well as TDNN + AAM [15] model. Improvement in the performance using relation networks increases as the number of speakers increases. Identifying the correct speaker within the enrollment set becomes more difficult with an increase in the number of speakers.

## 5. Conclusion

This paper proposed improved relation networks for speaker verification and identification. Additionally, we provided two enhancements to our technique - global classification for better supervision and an improved training regime for better utilization of samples present in an episode. The proposed method outperforms the existing techniques on speaker verification and identification on VoxCeleb, SITW and VCTK datasets.

## 6. 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, 2011.
- [2] J. Rohdin, A. Silnova, M. Diez, O. Plchot, P. Matějka, and L. Burget, "End-to-end dnn based speaker recognition inspired by i-vector and plda," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4874–4878.
- [3] S. J. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *2007 IEEE 11th International Conference on Computer Vision*, 2007, pp. 1–8.
- [4] D. Garcia-Romero and C. Y. Espy-Wilson, "Analysis of i-vector length normalization in speaker recognition systems," in *INTER-SPEECH*, 2011.
- [5] Z. Bai and X.-L. Zhang, "Speaker recognition based on deep learning: An overview," 2020. [Online]. Available: <https://arxiv.org/abs/2012.00931>
- [6] J.-w. Jung, H.-s. Heo, j.-h. Kim, H.-j. Shim, and H.-j. Yu, "Rawnet: Advanced end-to-end deep neural network using raw waveforms for text-independent speaker verification," *Proc. Interspeech 2019*, pp. 1268–1272, 2019.
- [7] H.-S. Lee, Y. Tso, Y.-F. Chang, H.-M. Wang, and S.-K. Jeng, "Speaker verification using kernel-based binary classifiers with binary operation derived features," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 1660–1664.
- [8] L. Ferrer, M. McLaren, and N. Brümmner, "A speaker verification backend with robust performance across conditions," *Comput. Speech Lang.*, vol. 71, no. C, jan 2022. [Online]. Available: <https://doi.org/10.1016/j.csl.2021.101258>
- [9] M. Kumar, T. Jin-Park, S. L. Bishop, and S. S. Narayanan, "Designing neural speaker embeddings with meta learning," *ArXiv*, vol. abs/2007.16196, 2020.
- [10] J. Wang, K.-C. Wang, M. T. Law, F. Rudzicz, and M. Brudno, "Centroid-based deep metric learning for speaker recognition," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 3652–3656.
- [11] T. Ko, Y. Chen, and Q. Li, "Prototypical networks for small footprint text-independent speaker verification," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6804–6808.
- [12] S. M. Kye, Y. Jung, H. Lee, S. J. Hwang, and H. Kim, "Meta-learning for short utterance speaker recognition with imbalance length pairs," in *INTERSPEECH*, 2020.
- [13] Y. Chen, W. Guo, and B. Gu, "Improved meta-learning training for speaker verification," in *INTERSPEECH*, 2021.
- [14] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, "Learning to compare: Relation network for few-shot learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [15] B. Desplanques, J. Thienpondt, and K. Demuyne, "Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," in *INTERSPEECH*, 2020.
- [16] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: A large-scale speaker identification dataset," *Interspeech 2017*, Aug 2017. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-950>
- [17] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *Interspeech 2018*, Sep 2018. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1929>
- [18] M. McLaren, L. Ferrer, D. Castán, and A. D. Lawson, "The speakers in the wild (sitw) speaker recognition database," in *INTER-SPEECH*, 2016.
- [19] C. Veaux, J. Yamagishi, and K. MacDonald, "Superseded - cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2016.
- [20] T. M. Hospedales, A. Antoniou, P. Micaelli, and A. J. Storkey, "Meta-learning in neural networks: A survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. PP, 2021.
- [21] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 4080–4090.
- [22] S.-H. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. Torr, "Res2net: A new multi-scale backbone architecture," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 2, p. 652–662, Feb 2021. [Online]. Available: <http://dx.doi.org/10.1109/TPAMI.2019.2938758>
- [23] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7132–7141.
- [24] D. Povey, A. K. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. K. Goel, M. Hannemann, P. Motlíček, Y. Qian, P. Schwarz, J. Silovský, G. Stemmer, and K. Veselý, "The kaldi speech recognition toolkit," 2011.
- [25] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *ArXiv*, vol. abs/1510.08484, 2015.
- [26] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5220–5224.
- [27] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *Interspeech 2019*, Sep 2019. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-2680>
- [28] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4685–4694.