



Domain Agnostic Few-shot Learning for Speaker Verification

Seunghan Yang, Debasmit Das, Janghoon Cho, Hyoungwoo Park, Sungrack Yun

Qualcomm AI Research[†]

{seunghan, debadas, janghoon, hwoopark, sungrack}@qti.qualcomm.com

Abstract

Deep learning models for verification systems often fail to generalize to new users and new environments, even though they learn highly discriminative features. To address this problem, we propose a few-shot domain generalization framework that learns to tackle distribution shift for new users and new domains. Our framework consists of domain-specific and domain-aggregation networks, which are the experts on specific and combined domains, respectively. By using these networks, we generate episodes that mimic the presence of both novel users and novel domains in the training phase to eventually produce better generalization. To save memory, we reduce the number of domain-specific networks by clustering similar domains together. Upon extensive evaluation on artificially generated noise domains, we can explicitly show generalization ability of our framework. In addition, we apply our proposed methods to the existing competitive architecture on the standard benchmark, which shows further performance improvements.

Index Terms: Few-shot Learning, Domain Generalization, Pseudo-label Clustering, Fine-tuning

1. Introduction

Deep learning models often fail to generalize to new domains and new classes, even though they produce highly discriminative representations. Domain generalization [1–4] and few-shot learning [5–7] have produced attractive solutions to address problems about unseen domains and classes, respectively. However, only few studies [8, 9] have focused on both problems, simultaneously. In particular, identification systems always face both problems in the test phase while most approaches only consider the same set of categories and domains as they are trained upon. Consequently, these closed-set models perform poorly on novel categories and novel domains. This paper addresses this problem for speaker verification, which is the task of accepting or rejecting the claimed identity of a speaker test utterance based on few enrolled utterances.

Previous approaches [10–18] on speaker verification train an embedding network using classification or metric learning loss on a dataset consisting of multiple speakers (users) in various domains. However, these methods do not generalize well to novel users and novel domains not involved during training. There are several works [19–21] using few-shot learning approaches that generalize to new users. In training, they exploit prototypical networks [6], and apply an episodic learning scheme to mimic the test conditions of new users. For classification, they use prototypes [19] or a class-centroid-based formulation [21] to compute the distance functions over all the classes. These meta-learning based methods consider generalization on novel users but do not consider novel domains. To achieve both objectives, we propose a domain generalization component that

mimics the presence of both novel users and novel domains in the episodic training phase.

In our framework, we consider two types of episode generations with domain-specific and domain-aggregation prototypical networks. Firstly, we generate the episodes consisting of the support and query set to mimic the cases of enrolling a novel user with few labeled samples (support set) and correctly classifying the test samples (query set). Here, the episodes are sampled from a specific domain and all domains for the domain-specific and domain-aggregation network, respectively. With the generated episodes and the prototypical loss function [6, 22], both networks are trained to adapt well to a new user. Secondly, we generate the episodes to mimic the cases of testing the network in a novel domain. Here, we consider a domain mismatched way: given a domain and its domain-specific network, the support and query sets are sampled from the other domain which is unseen to the specific network. Then, the support and query sets are fed into the domain-specific and domain-aggregation networks, respectively. With this episode generation and the prototypical loss function, we can train the domain-aggregation network with better domain generalization ability.

For the speaker verification task, the domains can be various environments such as background noises, recording devices, etc. However, modeling the domain-specific network for each domain may require much memory and computations. For the sake of efficiency, we consider a *pseudo* domain label approach where similar domains are clustered together using style features adopted in [23, 24].

We extensively evaluated our framework with two types of domain differences: (a) artificially generated and (b) those intrinsically present in a dataset. Firstly, using the VoxCeleb1 dataset [25], we artificially generated several domains with noise augmentations and split them into training and testing set. In the experiments, we observed that the proposed framework shows better generalization ability for both new speakers and new domains in test dataset. Also, the pseudo domain labels are beneficial for efficient training and further performance boost. Secondly, our framework when applied on top of [22], learns representations that can tackle intrinsic domain differences in VoxCeleb2 [26]. The results also show performance improvements on the standard evaluation benchmark: VoxCeleb1, SITW [27] and CNCeleb [28].

2. Proposed Framework

2.1. Task Description

In this paper, our task is speaker verification, which is a decision process that accepts or rejects an input utterance \mathbf{x} by comparing the utterance with a trained reference speaker model \mathbf{X}_{ref} . Mathematically, the decision process is represented as:

$$f(\mathbf{X}_{ref}, \mathbf{x}) \underset{reject}{\overset{accept}{\geq}} \tau \quad (1)$$

[†] Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

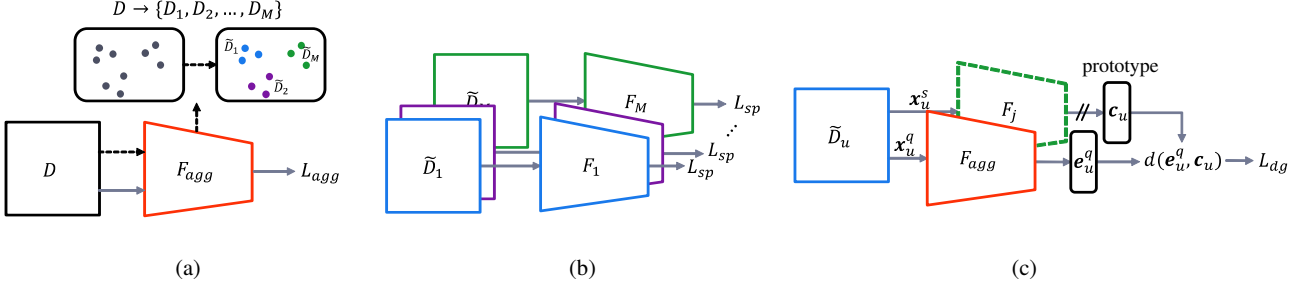


Figure 1: Our framework consists of: (a) Domain clustering: we assign pseudo domain labels to all samples using domain-discriminative features from the domain-aggregation network F_{agg} pre-trained by L_{agg} ; (b) Training domain-specific network: F_j is trained on the corresponding source domain after clustering and re-assignment of domains; (c) Domain-agnostic few-shot learning: episodes are generated to mimic the few-shot test case for novel users and domain-mismatched case for novel domains.

where $f(\cdot, \cdot)$ is a similarity metric score between \mathbf{X}_{ref} and \mathbf{x} . If the score is greater than a pre-defined threshold τ , \mathbf{x} is accepted as the utterance of the reference speaker; otherwise, \mathbf{x} is presumed to be from an impostor. Our goal is to train a speaker verification model that generalizes to novel users and domains.

2.2. Problem Setting and Notation

We have a source dataset $D = \{D_1, D_2, \dots, D_N\}$ consisting of multiple speakers from N domains, where D_j indicates the j^{th} domain. Each domain contains utterances and their corresponding speaker labels such that $D_j = \{(\mathbf{x}_j^i, y_j^i)\}_{i=1}^{n_j}$, where n_j is the number of samples in the j^{th} domain. We use the source data to train an embedding network F_{agg} such that it generalizes well to novel target users in novel testing domains D_* with different characteristics from the source domains. Note that it is different from data augmentation techniques that use noise and reverberation to enforce test domains to be included in the training domains. We denote a domain-specific network and a domain-aggregation network by F_j and F_{agg} , respectively. F_j specialized in j^{th} domain is learned so that F_{agg} learned with all domains can extract domain agnostic features even for novel domain data. Fig. 1 shows our overall framework.

2.3. Domain Clustering

Our learning scheme requires N domain-specific networks that consume lots of memory in training, which we can reduce by modifying the original source dataset D with pseudo domain labels by clustering. We extract domain-discriminative features and cluster them to assign the same pseudo domain labels to the samples of similar domains. For domain-discriminative features, we exploit the features proposed in style transfer algorithms for different image domains [23, 24, 29] that are recently shown applicable to audio domains [30–32]. In these methods, the mean and the standard deviation that are calculated across the spatial dimensions represent the input style. We obtain the domain-discriminative feature of an input \mathbf{x} by stacking these feature statistics as follows:

$$\{\mu(\phi_{agg,1}(\mathbf{x})), \sigma(\phi_{agg,1}(\mathbf{x})), \dots, \mu(\phi_{agg,l}(\mathbf{x})), \sigma(\phi_{agg,l}(\mathbf{x}))\}$$

where $\phi_{agg,l}$ indicates the output of the l^{th} layer in F_{agg} . Before extracting features, we first train F_{agg} with the source domain dataset D . Then, with the extracted domain-discriminative features, we perform k-means clustering [33] to obtain M clusters and assign the pseudo domain label to each sample. This is illustrated in Fig. 1(a). Accordingly, we obtain

a modified source dataset $D = \{\tilde{D}_1, \tilde{D}_2, \dots, \tilde{D}_M\}$ with pseudo domain labels, and $M \leq N$. In contrast to [24], our objective of assigning pseudo domain labels is to increase memory efficiency by clustering ambiguous domains and reducing the number of domain-specific networks.

2.4. Episodic Training

In this step, we describe our two types of episode generations to mimic the test cases of the novel user and novel domain. First, we split the dataset D into the support set $D^s = \{(\mathbf{x}^i, y^i)\}_{i=1}^{CK}$ and the query set D^q to mimic the few-shot test case for novel users, which is a C -way K -shot problem. Here, way and shot stand for the number of speakers and samples per speaker in each episode, respectively. We use the prototypical network [6] to be trained in an episodic manner such that it computes speaker prototypes with the support set and forces the query set to minimize the distance from the corresponding prototypes. Note that our algorithm can be adapted to any prototype-based loss functions, e.g., angular prototypical loss [22].

As illustrated in Fig. 1(b), the domain-specific network F_j is trained using episodes generated from \tilde{D}_j , and prototype of speaker k , \mathbf{c}_k , is obtained by averaging over embedding vectors in j^{th} domain support set:

$$\mathbf{c}_k = \frac{1}{K} \sum_{\{(\mathbf{x}_j^i, y_j^i)\} \in \tilde{D}_j^s} F_j(\mathbf{x}_j^i) \mathbb{1}(y_j^i = k) \quad (2)$$

where $\mathbb{1}(\cdot)$ is the indicator function returning 1 for true statements and 0 otherwise. The episodic training loss L_{sp} for the domain-specific network is defined as follows:

$$L_{sp} = \sum_{\{(\mathbf{x}_j^i, y_j^i)\} \in \tilde{D}_j^q} -\log(p(y = y_j^i | \mathbf{x}_j^i)) \quad (3)$$

$$\text{where } p(y = y_j^i | \mathbf{x}_j^i) = \frac{\exp(-d(F_j(\mathbf{x}_j^i), \mathbf{c}_{y_j^i}))}{\sum_{k'} \exp(-d(F_j(\mathbf{x}_j^i), \mathbf{c}_{k'}))},$$

and $d(\cdot)$ is the Euclidean distance. Similarly, the domain-aggregation network F_{agg} is trained by the loss function L_{agg} , which is the same with L_{sp} except that the support sets and the query sets are sampled from all the source domains D .

The second type of episode generation and training is illustrated in Fig. 1(c). Our objective is to train F_{agg} to extract discriminative features even for the novel domain data. Hence, we construct the support set \tilde{D}_u^s and query set \tilde{D}_u^q from u^{th} domain and mimic the test case of domain generalization exploiting F_j

(domain-mismatch). With the domain-mismatch condition, *i.e.*, $j \neq u$, unseen domain prototypes \mathbf{c}_k are calculated by Eq. 2 using F_j on the novel domain set \tilde{D}_u^s , and embedding vectors of \tilde{D}_u^q are extracted from F_{agg} . The loss L_{dg} for few-shot domain generalization episodes is as follows:

$$L_{dg} = \sum_{\{(\mathbf{x}_u^i, y_u^i)\} \in \tilde{D}_u^q} -\log(p(y = y_u^i | \mathbf{x}_u^i)) \quad (4)$$

where
$$p(y = y_u^i | \mathbf{x}_u^i) = \frac{\exp(-d(F_{agg}(\mathbf{x}_u^i), \mathbf{c}_{y_u^i}))}{\sum_{k'} \exp(-d(F_{agg}(\mathbf{x}_u^i), \mathbf{c}_{k'}))}.$$

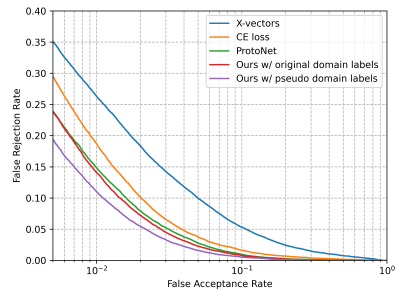
Since the u^{th} domain data is novel to F_j , F_{agg} learns to map embedding vectors in a domain-agnostic way. Here, only F_{agg} is updated by L_{dg} since we need to keep \tilde{D}_u novel to F_j during training. Our training procedure is summarized as follows: 1) Train F_{agg} using L_{agg} and extract domain-discriminative features, 2) Perform clustering on extracted features and assign pseudo domain labels: this leads to M domains from N input domains, 3) Generate M domain-specific networks and train them using Eq. 3, 4) Simultaneously, train F_{agg} using the combined loss $L_{agg} + \lambda_{dg}L_{dg}$ where λ_{dg} is a balancing parameter.

3. Experimental Results

3.1. Experiment Setup

Dataset. First, to analyze generalization ability of our framework, we designed the evaluation experiment of simulating domain differences by artificially producing data augmentations. We selected 1,088 speakers in VoxCeleb1 [25], consisting of 1,000 speakers for training and 88 for testing. Each speaker has 45 utterances, and each utterance was randomly cropped into 2-sec-long for the use-cases of short utterance enrollment and testing. To explicitly generate several domains, we augmented domains by mixing noises. Then, we split them into train and test sets to simulate the test condition where data are recorded in novel domains. The source dataset consists of 4 domains: original set (clean) and three noisy domains which were augmented by synthesizing the clean with babble, car, and music noises (0dB SNR). Each of the novel speaker’s data in the target domain consists of 5 utterances for enrollment and 35 utterances for testing. The target domain consists of 7 domains: original set (clean) and 6 noisy domains which were augmented by synthesizing the clean with babble, car, music, office, ambient room, and typing noises (0dB SNR). Thus, our model is evaluated on 4 in-domain sets (*i.e.* clean, babble, car, music) and 3 out-domain sets (*i.e.* office, ambient room, and typing). Additionally, to analyze models in diverse environments, we augmented VoxCeleb1 with various noise types and SNR values for source and target domains. Details are described in Section 3.3.

Next, we evaluated our proposed methods on the standard benchmark following [22], where VoxCeleb2 [26] is used for training and VoxCeleb1 [25], SITW [27], and CNCeleb [28] are used for testing. Since VoxCeleb2 contains speech from speakers spanning a wide range of different ethnicities, accents, professions, and ages recorded in various environments, we assume that there exists domain difference intrinsically within the dataset. The development set of VoxCeleb2 contains 5,994 speakers in total and is entirely disjoint from the test datasets. VoxCeleb1 test set and SITW-Eval.Core are widely-used evaluation datasets, and it is worth noting that the acoustic condition of them is similar with that of the training set VoxCeleb2, but there is domain difference intrinsically present between and



Model	EER (%)	FRR @ FAR 10%
X-vectors	5.685	3.817
CE loss	3.879	1.716
ProtoNet	3.496	1.257
Ours w/ original domain labels	3.231	1.016
Ours w/ pseudo domain labels	2.717	0.756

Figure 2: ROC curve, EER, and FRR for different methods.

within the datasets. CNCeleb is a large-scale speaker dataset comprised of 11 different genres from 3,000 Chinese celebrities, and we utilize CNCeleb.E for testing. The acoustic condition of CNCeleb is severely different from that of VoxCeleb.

Implementation details. We use the x-vector system [13] and the speaker embedding network [18] for evaluating on the artificially generated dataset. The x-vector system is a standard DNN speaker verification system using temporal statistics. The speaker embedding network (SE network) uses an 1D-CNN architecture, and [18] shows that SE network can effectively learn with short utterances. We use both systems for the baseline, and we adapt our episodic learning algorithm to SE network. We trained the network with 5-way 5-shot training episodes where 5 support and 5 query points per speaker are included. Given 4 source domains (clean, babble, car, and music), we trained the domain-specific networks with two different setups: $N = 4$ (original) and $M = 3$ (clustered). We chose the balancing parameters as $\lambda_{dg} = 0.8$ and $\lambda_n = \{0.1, 0.2, 0.3\}$.

For evaluating our framework on the standard benchmark, we exploit Fast ResNet-34 with angular prototypical loss function [22] as a baseline. No data augmentation is performed during training, apart from the random sampling. We trained the model with 200-way 2-shot training following [22]. We adapt our episodic learning algorithm to Fast ResNet-34 with $\lambda_{dg} = 0.1$ using the same training strategy with the baseline for a fair comparison. We adapt our clustering algorithm to assign pseudo domain labels to training samples and use them for training domain-specific networks and conduct ablation studies with various number of domains $M = 2, 3, 4$.

Evaluation metrics. We used False Acceptance Rate (FAR) and False Rejection Rate (FRR). FAR is the percentage of identification instances in which unauthorised persons (imposters) are incorrectly accepted, while FRR is that in which the authorised person (target user) is incorrectly denied. The Equal Error Rate (EER) is measured at which FAR and FRR are the same. The minimum detection cost of the function (MinDCF) was defined by the NIST SRE evaluation plans [34]. DCF is a weighted sum of FRR and FAR as $C_{fr} \cdot FRR \cdot P_{target} + C_{fa} \cdot FAR \cdot (1 - P_{target})$. In particular, the parameters $C_{fr} = 1$, $C_{fa} = 1$ and $P_{target} = 0.05$ are used for the cost function.

3.2. Comparison Studies on Generated Domains

On the artificially generated domains, we compare our framework with the following methods: the x-vector system (X-vectors) and SE network trained with cross-entropy loss (CE loss) which do not consider both novel users and domains, the




Figure 3: *T-SNE* plot of the features obtained by ProtoNet (left) and our framework (right). Color and shape represent speaker and domains, respectively. *O* and *X* indicate the feature of in-domain and out-domain, respectively.

Table 1: Average EER (%) of four methods given in-domain and out-domain data with different numbers of domains.

Model	In-4	Out-3	Out-11	Out-6
CE loss	3.683	4.400	4.781	10.708
ProtoNet	2.888	3.649	4.055	9.707
Ours w/ original domain labels	2.895	3.417	3.873	9.265
Ours w/ pseudo domain labels	2.528	2.916	3.396	7.185

prototypical network (ProtoNet) which considers test case of novel users but not novel domains.

Fig. 2 shows the ROC curve and the table of EER and FRR of the 5 different methods. X-vectors are robust speaker embeddings given long utterances, but under short-utterance setting, where all utterances are cropped into 2-sec-long, SE network outperforms x-vectors. For this reason, we applied our learning scheme to the SE network and showed 4 different results in Fig. 2. Our method achieved the best performance compared with other learning frameworks. It shows that our learning scheme can lead the network to extract discriminative features even for the unseen users and domains. Surprisingly, our method w/ pseudo domain labels could achieve memory efficiency and produce less error compared with the original domain labels. For better domain generalization, the domain-specific network should have distinct domain characteristics, *i.e.*, trained with well-separated domain features. In our case, we define noise type as the domain, and some noise types could have similar characteristics in the feature space. Thus, pseudo domain labels obtained with clustering performed better than the original domain labels. Also, this leads to the reduction in memory requirements for domain-specific networks. At 10% FAR, ours w/ pseudo domain labels shows 55.94%, 39.86%, and 25.59% relative FRR improvements compared to the other three methods using SE network.

3.3. Additional Analyses

Visualization. As shown in the t-SNE plot [35] of embedding vectors in Fig. 3, our few-shot domain generalization framework extracts well-clustered features per each user in a domain-agnostic way. Features extracted from ProtoNet are clustered well per each user but some features of out-domains are far from those of in-domains and the cluster center. The features could be easily rejected during verification, and it leads to higher FRR.

Performance on in-domains and out-domains. In Table 1, the first and second column respectively indicates the EER given the test data from in-domain (4 source domains) and out-domain (3 novel domains) data. CE loss and ProtoNet use data augmentation to be robust to noise environments. However, they are vulnerable to novel domains that cannot be covered by noise augmentations. Our method could adapt well to out-domains by explicitly considering the domain generalization test case during training.

Table 2: Effect of number of pseudo domains on EER.

	Original domain labels	Pseudo domain labels			
	$N=10$	$M=10$	$M=8$	$M=6$	$M=4$
EER (%)	2.264	2.044	2.145	2.176	2.332

Table 3: EER and MinDCF of the baseline and ours with different number of pseudo domain labels (* is the number in [22]).

Model	VoxCeleb1		SITW		CNCeleb.E	
	EER	MinDCF	EER	MinDCF	EER	MinDCF
Fast ResNet-34 [22]	2.29 (*2.37)	0.188	4.02	0.301	15.83	0.815
Ours w/ $M=2$	2.09	0.177	3.94	0.278	15.50	0.717
Ours w/ $M=3$	2.12	0.174	3.91	0.285	15.44	0.735
Ours w/ $M=4$	2.16	0.177	3.88	0.280	16.04	0.727

Test on various target domains. We evaluated our framework using various out-domain settings: in Table 1, the third and the fourth column shows EER with 11 novel noise types (Out-11) and 3 novel noise types with an additional severe SNR value (Out-6). Especially, in Out-6 case, we applied 2 different SNR values (-6, 0dB) to 3 novel noise types. Our framework shows better results especially for the Out-6 which has large domain gaps due to severe noise. It demonstrates that the training scheme to deal with test case of unseen domains can lead the network to extract discriminative features even under the situation where train and test domains are severely different.

Train with various source domains. To study the effect of pseudo domain labels, we tested our framework with more source domains by augmenting original domains with 3 different SNR values (-6, 0, 6dB SNR). Overall, we have 10 domains: clean and 9 noisy domains. Training 10 domain-specific networks requires much memory and compute, and we combine the original domains by clustering. In Table 2, ours w/ pseudo domain labels with $M = 4$ that uses 4 domain-specific components achieved comparable performance compared to ours w/ original domain labels. In addition, ours w/ pseudo domain labels with $M = 10$ outperformed over ours w/ original domain labels, ensuring the efficacy of pseudo domain labels.

3.4. Evaluation on the Standard Benchmark

As shown in Table 3, our episodic learning strategy boosts the overall performance on all test datasets by a clear margin. Our network learns the generalization ability from environments intrinsically present within VoxCeleb2. Hence, it works better for novel environments where the domain difference between training and test sets is not large. *i.e.*, VoxCeleb1 and SITW. However, it is hard to improve the generalization ability on novel genres because VoxCeleb2 does not contain multiple genres, which indicates that our episodic learning setup cannot mimic the test scenario. Thus, the performance marginally increases on CNCeleb, where the data are recorded for various genres, *e.g.*, singing, movie, drama, etc. We expect that our episodic learning strategy on the training set of CNCeleb can improve the performance on these novel genres, which we leave for future experimental work.

4. Conclusion

We propose a domain generalization framework by generating two types of episodes to learn a speaker verification model that generalize to novel users and domains. We include domain clustering followed by learning domain-specific and -aggregation networks in source domain. The domain-aggregation network is learned to be domain-agnostic with episodes mimicking the test case of novel users and domains. Extensive experiments show that our framework produces better domain-agnostic features and considerable performance improvements.

5. References

- [1] H. Li, S. J. Pan, S. Wang, and A. C. Kot, "Domain generalization with adversarial feature learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [2] D. Li, Y. Yang, Y.-Z. Song, and T. Hospedales, "Learning to generalize: Meta-learning for domain generalization," in *Proceedings of the Conference on Artificial Intelligence (AAAI)*, 2018.
- [3] D. Li, J. Zhang, Y. Yang, C. Liu, Y.-Z. Song, and T. M. Hospedales, "Episodic training for domain generalization," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [4] J. Kang, R. Liu, L. Li, Y. Cai, D. Wang, and T. F. Zheng, "Domain-invariant speaker vector projection by model-agnostic meta-learning," in *Proceedings of the INTERSPEECH*, 2020.
- [5] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International Conference on Machine Learning (ICML)*, 2017.
- [6] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [7] D. Das, S. Yun, and F. Porikli, "ConFeSS: A framework for single source cross-domain few-shot learning," in *International Conference on Learning Representations (ICLR)*, 2022.
- [8] Y. Guo, N. C. Codella, L. Karlinsky, J. V. Codella, J. R. Smith, K. Saenko, T. Rosing, and R. Feris, "A broader study of cross-domain few-shot learning," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [9] H.-Y. Tseng, H.-Y. Lee, J.-B. Huang, and M.-H. Yang, "Cross-domain few-shot classification via learned feature-wise transformation," in *International Conference on Learning Representations (ICLR)*, 2020.
- [10] G. Heigold, I. Moreno, S. Bengio, and N. Shazeer, "End-to-end text-dependent speaker verification," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016.
- [11] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification," in *Proceedings of the INTERSPEECH*, 2017.
- [12] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018.
- [13] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [14] Y. Li, F. Gao, Z. Ou, and J. Sun, "Angular softmax loss for end-to-end speaker verification," in *International Symposium on Chinese Spoken Language Processing (ISCSLP)*, 2018.
- [15] H.-S. Heo, J.-w. Jung, I.-H. Yang, S.-H. Yoon, H.-j. Shim, and H.-J. Yu, "End-to-end losses based on speaker basis vectors and all-speaker hard negative mining for speaker verification," in *Proceedings of the INTERSPEECH*, 2019.
- [16] Y. Liu, L. He, and J. Liu, "Large margin softmax loss for speaker verification," in *Proceedings of the INTERSPEECH*, 2019.
- [17] Z. Ren, Z. Chen, and S. Xu, "Triplet based embedding distance and similarity learning for text-independent speaker verification," in *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2019.
- [18] S. Yun, J. Cho, J. Eum, W. Chang, and K. Hwang, "An end-to-end text-independent speaker verification framework with a keyword adversarial network," in *Proceedings of the INTERSPEECH*, 2019.
- [19] P. Anand, A. K. Singh, S. Srivastava, and B. Lall, "Few shot speaker recognition using deep neural networks," *arXiv preprint arXiv:1904.08775*, 2019.
- [20] 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 *Proceedings of the INTERSPEECH*, 2020.
- [21] J. Wang, K.-C. Wang, M. T. Law, F. Rudzicz, and M. Brudno, "Centroid-based deep metric learning for speaker recognition," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [22] 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 *Proceedings of the INTERSPEECH*, 2020.
- [23] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [24] T. Matsuura and T. Harada, "Domain generalization using a mixture of multiple latent domains," in *Proceedings of the Conference on Artificial Intelligence (AAAI)*, 2020.
- [25] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," in *Proceedings of the INTERSPEECH*, 2017.
- [26] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *Proceedings of the INTERSPEECH*, 2018.
- [27] M. McLaren, L. Ferrer, D. Castan, and A. Lawson, "The speakers in the wild (sitw) speaker recognition database," in *Proceedings of the INTERSPEECH*, 2016, pp. 818–822.
- [28] L. Li, R. Liu, J. Kang, Y. Fan, H. Cui, Y. Cai, R. Vippera, T. F. Zheng, and D. Wang, "Cn-celeb: multi-genre speaker recognition," *Speech Communication*, 2022.
- [29] D. Jung, S. Yang, J. Choi, and C. Kim, "Arbitrary style transfer using graph instance normalization," in *Proceedings of the IEEE international conference on image processing (ICIP)*, 2020.
- [30] B. Kim, S. Yang, J. Kim, and S. Chang, "Domain generalization on efficient acoustic scene classification using residual normalization," *arXiv preprint arXiv:2111.06531*, 2021.
- [31] B. Kim, S. Yang, J. Kim, H. Park, J.-T. Lee, and S. Chang, "Towards robust domain generalization in 2d neural audio processing," 2021.
- [32] B. Kim, S. Yang, J. Kim, and S. Chang, "Qti submission to dcase 2021: Residual normalization for device imbalanced acoustic scene classification with efficient design," *DCASE2021 Challenge, Tech. Rep*, 2021.
- [33] J. MacQueen *et al.*, "Some methods for classification and analysis of multivariate observations," in *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, 1967.
- [34] O. Sadjadi, C. Greenberg, E. Singer, L. Mason, and D. Reynolds, "Nist 2021 speaker recognition evaluation plan," 2021-07-12 04:07:00 2021.
- [35] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.