Meta-learning with Latent Space Clustering in Generative Adversarial Network for Speaker Diarization

Monisankha Pal, *Member, IEEE,* Manoj Kumar, *Member, IEEE,* Raghuveer Peri, *Member, IEEE,* Tae Jin Park, *Member, IEEE,* So Hyun Kim, Catherine Lord, Somer Bishop, and Shrikanth Narayanan, *Fellow, IEEE*

Abstract—The performance of most speaker diarization systems with x-vector embeddings is both vulnerable to noisy environments and lacks domain robustness. Earlier work on speaker diarization using generative adversarial network (GAN) with an encoder network (ClusterGAN) to project input x-vectors into a latent space has shown promising performance on meeting data. In this paper, we extend the ClusterGAN network to improve diarization robustness and enable rapid generalization across various challenging domains. To this end, we fetch the pre-trained encoder from the ClusterGAN and fine-tune it by using prototypical loss (meta-ClusterGAN or MCGAN) under the meta-learning paradigm. Experiments are conducted on CALL-HOME telephonic conversations, AMI meeting data, DIHARD II (dev set) which includes challenging multi-domain corpus, and two child-clinician interaction corpora (ADOS, BOSCC) related to the autism spectrum disorder domain. Extensive analyses of the experimental data are done to investigate the effectiveness of the proposed ClusterGAN and MCGAN embeddings over x-vectors. The results show that the proposed embeddings with normalized maximum eigengap spectral clustering (NME-SC) back-end consistently outperform Kaldi state-of-the-art x-vector diarization system. Finally, we employ embedding fusion with x-vectors to provide further improvement in diarization performance. We achieve a relative diarization error rate (DER) improvement of 6.67% to 53.93% on the aforementioned datasets using the proposed fused embeddings over x-vectors. Besides, the MCGAN embeddings provide better performance in the number of speakers estimation and short speech segment diarization as compared to x-vectors and ClusterGAN in telephonic data.

Index Terms—ClusterGAN, MCGAN, NME-SC, speaker diarization, speaker embeddings, x-vector.

I. INTRODUCTION

Speaker diarization [1], the task of determining "who spoke when" in a multi-speaker audio stream has a wide range of applications such as information retrieval, speaker-based indexing, meeting annotations, and conversation analysis [2]. Present-day diarization systems typically comprise four components: (a) A speech segmentation module that removes the non-speech parts using a speech activity detector (SAD) and segments the speech part into multiple speaker-homogeneous short segments [3]; (b) A speaker representation

M. Pal, M. Kumar, R. Peri, T. J. Park and S. Narayanan are with Signal Analysis and Interpretation Laboratory, University of Southern California, Los Angeles, USA e-mail: (mp_323@usc.edu; prabakar@usc.edu; rperi@usc.edu; taejinpa@usc.edu; shri@sipi.usc.edu).

S. H. Kim is with Center for Autism and the Developing Brain, Weill Cornell Medicine, USA e-mail:(sok2015@med.cornell.edu)

C. Lord is with Semel Institute of Neuroscience and Human Behavior, University of California Los Angeles, USA e-mail:(CLord@mednet.ucla.edu)

S. Bishop is with Department of Psychiatry, University of California, San Francisco, USA e-mail:(somer.bishop@ucsf.edu)

(embedding) extractor that maps the segments into fixed-dimensional *speaker embeddings* such as i-vectors [4], [5], d-vectors [6], [7] and *x-vectors* [3], [8]; (c) A clustering module that determines the number of constituent speakers in an audio recording and clusters the extracted embeddings into these speakers [9], [10]; (d) A re-segmentation module that refines the clustering results [3].

For embedding extraction, typically i-vectors have been obtained through total variability space projection [11]. However, recently significant performance improvement has been shown using deep neural network embeddings such as d-vectors with architectures such as LSTM [6], [12], CNN [13]; and x-vectors with time-delay neural network (TDNN) [3], [14]. The combination of different embeddings, e.g., c-vectors using 2D self-attentive structure, has also been proposed to exploit the complementary merits of each embedding [15].

In terms of clustering, most of the existing algorithms that have been used in speaker diarization are unsupervised. Among them, agglomerative hierarchical clustering (AHC) [3] and spectral clustering (SC) [16] using pairwise embedding similarity measurement techniques like cosine distance [6], [9], PLDA [17] and using an LSTM [10] are the most popular. Similarly, other unsupervised clustering methods such as Gaussian mixture model [4], [13], mean-shift [5], k-means [18], and links [19] have also been adopted for speaker diarization. Moreover, clustering depends on tuning hyperparameters like stopping threshold (for AHC), the p-value for binarization of affinity matrix (for SC). However, more recently, an auto-tuning and improved version of the spectral clustering approach on x-vectors using cosine similarity measure, which is called as *normalized maximum eigengap spectral clustering (NME-SC)* was introduced in [9]. Despite the success of these speaker clustering algorithms, speaker diarization remains a challenging task due to the wide heterogeneity and variability of audio data recorded in many real-world scenarios [20].

The other approach for speaker clustering has been based on supervised methods. A fully supervised speaker diarization framework, named UIS-RNN was proposed in [7]. Although this model for clustering produces excellent performance in telephone conversations, its performance deteriorates in a more challenging multi-domain database like DIHARD II [21]. To improve the UIS-RNN diarization performance further, a novel sample-mean loss function to train the RNN has been introduced very recently [21]. Efforts have been made to automatically deal with speaker-overlapping speech segments

and directly optimize an end-to-end neural network based on diarization errors [22]. The network is trained in a supervised manner using a permutation free objective function. The diarization performance was further enhanced by introducing a self-attention based end-to-end neural network [23]. Although the above methods do not rely on clustering and can directly compute the final diarization outputs using a single network, they assume that the number of speakers is known apriori or at least bounded to two speakers. Along these lines, the performance of deep embedded clustering, which was originally proposed in [24], was incorporated and modified for speaker clustering in diarization task [18]. The limitation of this work is that a good estimate of the number of speakers is needed for its evaluation.

While performance of tasks such as speech and speaker recognition have improved significantly due to supervised deep learning approaches, most of the speaker clustering is yet to take advantage of similar techniques. The main problem that hinders in making clustering a supervised task is associated with the fact that speaker labels are ambiguous (e.g., both "112233" and "223311" sequences of labels are equally correct for the same diarization session). In our earlier proposed work, we incorporated *ClusterGAN* to non-linearly transform DNN-based speaker embeddings into a low-dimensional latent space better suited for clustering [25]. The proposed ClusterGAN, which exploits the GAN latent space with the help of an encoder network, was trained with a combination of adversarial loss, latent variable recovery loss, and clustering-specific loss. Although the proposed system showed significant performance improvement over x-vector based state-of-the-art in meeting and child-adult interaction corpora, its performance was not tested against telephone conversations and a broader set of multi-domain data.

In this work, a ClusterGAN network which was originally proposed for image clustering [26], is adopted and modified for the speaker clustering task in the speaker diarization framework. The GAN and the encoder network are trained jointly in a supervised manner with clustering-specific loss and latent embeddings are extracted using the trained encoder to perform unsupervised clustering at the back-end. Two main advantages of GAN-based latent space clustering are the interpretability and interpolation in the latent space [26]. We use ClusterGAN-trained encoder network as initialization to further fine-tune it with meta-learning based *prototypical loss* function [27], [28]. This is represented as *meta-ClusterGAN or MCGAN* in this paper. The prototypical network was introduced for the few-shot image classification task [27] and is the state-of-the-art approach on a few-shot image classification benchmark. The motivation behind using proto-learning for our task is that it has a simpler inductive bias in the form of speaker prototypes and can perform rapid generalization to new speakers or types of data not seen while training. The prototypical loss trained for learning a metric space to mimic the test scenario will be beneficial in capturing information related to both generalization and clustering objectives.

The main contributions of this paper are: (a) A novel speaker diarization framework based on prototypical learning; (b) Extensive multi-domain experimental evaluation and analysis

of the proposed diarization system on various challenging speaker diarization corpora; (c) Demonstration of the use of novel speaker embeddings that outperform x-vectors through analysis across various challenging scenarios.

II. RELATED WORKS

A. Deep clustering algorithms

Using deep neural networks to non-linearly transform the input data into cluster-friendly representation along with dimension reduction is commonly known as deep clustering [29]. Recent deep clustering methods on image data using autoencoder networks like deep embedded clustering (DEC) [24] achieve impressive clustering performance. Generative modeling based approaches like variational deep embedding [30], information maximizing GAN (InfoGAN) [31], GAN mixture model [32] learn latent representation space and can interpolate to generate new samples from the data distribution. In all these algorithms, the deep neural network is usually trained on two types of losses: representation loss or network loss and clustering-specific loss. The network loss is essential for network initialization and is used to learn feasible latent features. The different network losses are reconstruction loss of autoencoder, variational loss of a variational autoencoder, and adversarial loss of GANs. On the other hand, clustering-specific loss helps to learn representations suitable for clustering. The option for clustering-specific losses are assignment losses like k-means loss [33], cluster assignment hardening loss [24], agglomerative clustering loss [34], spectral clustering loss [35] or regularization losses such as locality preserving loss, cluster classification loss [29]. Different ClusterGAN proposed for image data clustering adopts adversarial loss in GAN and clustering-specific loss like balanced self-paced entropy minimization loss [36] or cluster classification loss [26]. Very recently, few deep clustering approaches like transformer-based discriminative neural clustering model [37], deep clustering loss in end-to-end neural speaker diarization [23], deep embedded clustering [18], and ClusterGAN [25] have been used for speaker diarization. Although multifarious deep clustering approaches have been successfully applied for image data clustering, their application toward speaker diarization has been limited mainly due to the problem of the unknown number of speakers in a given diarization session.

B. Meta-learning algorithms

Inspired by human learning of new categories (classes) given just a very few examples, the meta-learning model trained over a large variety of learning tasks can adapt or generalize well to potentially unseen tasks [38]. It is also known as learning-to-learn, which learns on a given task and also across tasks. In the computer vision literature, there are three common approaches to meta-learning: metric-based [39], model-based [40], and optimization-based [41]. Metric learning aims at learning a metric or distance function over the embedding space. Among metric-learning based approaches, Siamese networks [42] and triplet networks [43] for learning speaker embeddings have been proposed for speaker recognition [44] and speaker diarization [45], and have yielded promising performances. The prototypical network that learns a metric space by computing prototype representation of
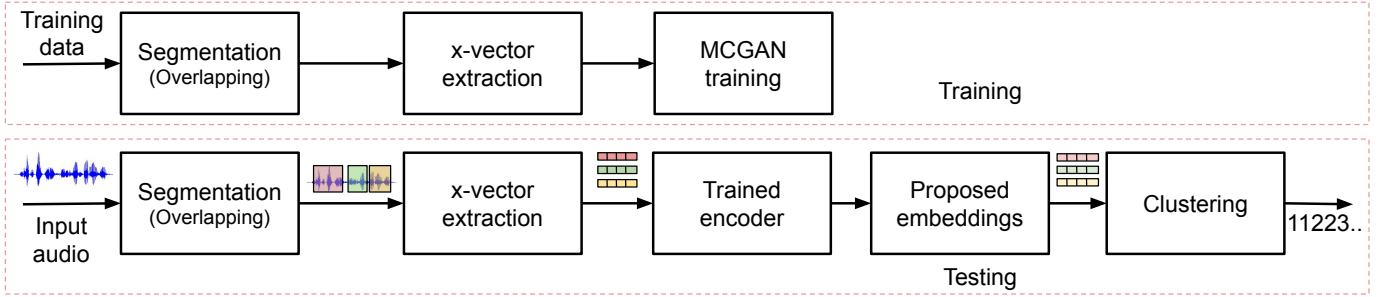
Fig. 1. Skeleton of the proposed speaker diarization system.

each class is a state-of-the-art approach for a few-shot image classification tasks [27]. Along these lines, prototypical loss to optimize a speaker embedding model for the speaker verification task was explored in [28], [44]. The resulting model provides superior performance to triplet loss based models. Very recently, the usage of protonets for child-adult audio classification task was explored in [46]. Our proposed approach uses prototypical loss (PTL) to fine-tune the encoder of ClusterGAN for robust speaker embedding extraction in the speaker diarization framework.

III. PROPOSED SPEAKER DIARIZATION SYSTEM

An overview of our proposed speaker diarization system is shown in Fig. 1. The non-speech part in a given multi-speaker conversation is removed first by using a speech activity detection (SAD) system. Our diarization system uses Kaldi[1] style uniform segmentation and the segments are embedded into a fixed-dimensional vector using a time-delay neural network (TDNN), which is commonly known as x-vector [14]. The proposed meta-ClusterGAN (MCGAN) is developed on top of x-vectors to perform deep latent space clustering for speaker diarization. We describe each of the modules in the diarization pipeline below.

A. Segmentation

In this paper, our proposed system uses oracle SAD for all the analysis and experiments, following common practice in the speaker diarization literature [3], [7], [47]. Therefore, our approach starts with a temporal uniform segmentation of 1.5 sec with an overlap of 1 sec between two adjacent segments. This denser segmentation gives more number of samples while evaluating a diarization session and it helps in clustering.

B. Speaker embedding vector

The speaker embedding vectors used to train the Cluster-GAN models are x-vectors, which are fixed-length representation using a TDNN from variable-length utterances. In this approach, MFCCs are first extracted at frame-level and input to a TDNN for supervised training using the categorical cross-entropy loss based on the speaker labels. The statistics pooling layer inside the TDNN architecture is used to convert frame-level features into a segment-level embedding. The detailed procedure of x-vector extraction is concisely described in [3], [14]. In this paper, we use Kaldi-based pre-trained x-vectors.

[1]https://kaldi-asr.org/

C. Meta-ClusterGAN (MCGAN) training

We employ meta-ClusterGAN (MCGAN) for the speaker embedding extraction in the speaker diarization framework. The motivation behind introducing MCGAN is to non-linearly transform the input x-vectors (trained with categorical cross-entropy loss) into another embedding suitable for speaker clustering and that can generalize well to new classes (here, speakers) not seen while training. As shown in Fig. 2, the proposed MCGAN training has two phases: (a) parameter initialization using a ClusterGAN, trained with clustering-specific loss in GAN latent space, and (b) inducing robustness to the initialized encoder in ClusterGAN by further fine-tuning it with meta-learning based prototypical loss. We discuss each phase of training further in the following subsections.

1) ClusterGAN training: We pre-train the MCGAN encoder using ClusterGAN training since it can decipher the original data representation by exploiting the GAN latent space. The learned encoder in ClusterGAN can generate embeddings in another space while maintaining the separable properties among the classes. ClusterGAN comprises three components: generator (G), discriminator (D) and encoder (E). The complete ClusterGAN architecture details and its training procedure are described in detail below.

a) Adversarial training: ClusterGAN adopts adversarial training of GANs for the clustering task. The standard GAN is formulated as an adversarial mini-max game between two neural networks: a generator (G) and a discriminator (D) [48]. The generator aims to create a map from latent space to data space, i.e., $G : \mathbf{z} \to \hat{\mathbf{x}}$. It takes random noise $\mathbf{z}$ sampled from $p_{\mathbf{z}}$ and synthesizes data similar to original data to fool the discriminator. The discriminator is considered to be a mapping from the data space to a real value $D : \mathbf{x} \to \mathbb{R}$. It takes real data $\mathbf{x}$ sampled from $p_{\mathbf{x}}^r$ and aims to distinguish the real data from the generator produced samples. Although GANs can learn to mimic any data distribution, they are difficult to train due to the mode collapse problem [49]. To address this issue, several variants of GANs such as Wasserstein GAN (WGAN) [49], and improved WGAN [50] (IWGAN) have been proposed in the literature. In this work, we incorporate IWGAN as our GAN network. The objective function of this adversarial game between G and D is

$$\min_{G} \max_{D} U_{\text{IWGAN}}(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}^r} [D(\mathbf{x})] - \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}} [D(G(\mathbf{z}))] + \lambda \cdot \text{GP}$$

(1)

where λ denotes the gradient penalty coefficient and GP represents the gradient penalty term [50]. The gradient penalty
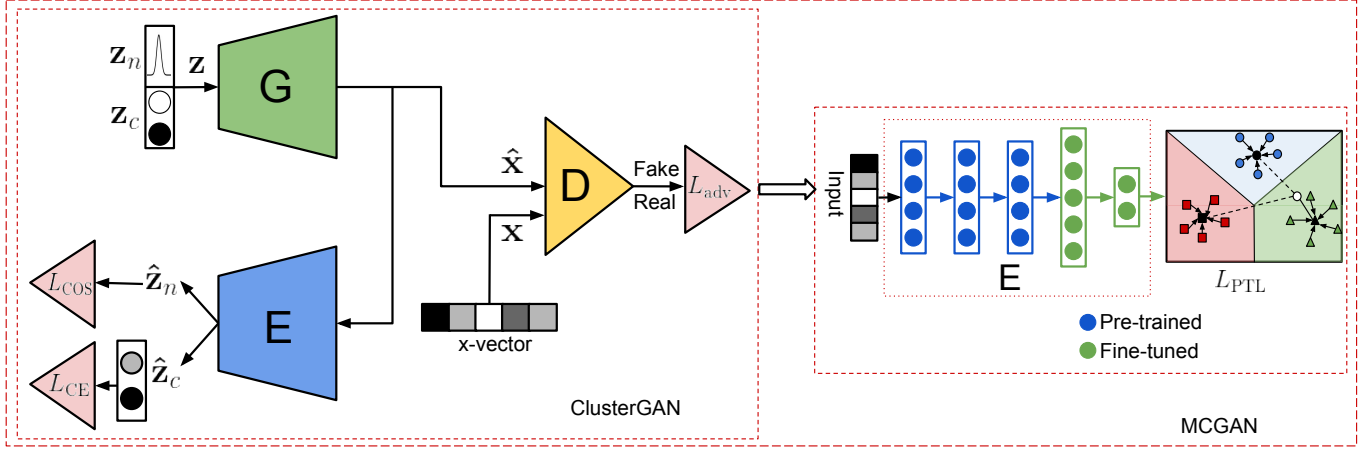
Fig. 2. MCGAN architecture. Here, L_{adv}, L_{COS}, L_{CE} and L_{PTL} represent adversarial, cosine distance, cross-entropy and prototypical loss functions.

term can be expressed as

$$\text{GP} = \mathbb{E}_{\hat{\mathbf{x}} \sim p_{\hat{\mathbf{x}}}} \left[\left(\|\nabla_{\hat{\mathbf{x}}} D(\hat{\mathbf{x}})\|_2 - 1 \right)^2 \right] \qquad (2)$$

where $\hat{\mathbf{x}} = \epsilon\mathbf{x} + (1 - \epsilon)G(\mathbf{z})$ and ϵ is a random number uniformly sampled in between 0 and 1.

b) Mixture of discrete and continuous latent variables: One possible way to perform clustering in the latent space is to back-project the data into the GAN latent space and then cluster it. The latent vectors for GANs trained with different priors such as Gaussian or uniform distribution usually lead to bad clustering [51]. Although the latent space may contain useful information about the data, the distance geometry does not reflect any form of clustering. To combat this issue, boosting the latent space using categorical variables ($\mathbf{z}_c$) to form non-smooth geometry is essential. The discrete variable $\mathbf{z}_c$ as a mixture with the continuous random variable ($\mathbf{z}_n$) will restrict the GAN generator to produce each mode only generating samples from a corresponding category in the real data. A similar type of latent variable structure within a GAN generator for learning disentangled and meaningful representation was employed in InfoGAN [31]. However, ClusterGAN has been reported to be superior to InfoGAN for clustering [26]. Furthermore, continuity in the latent space is also required for good interpolation objective and GANs have good interpolation ability. Therefore, our latent variable $\mathbf{z}$ is a concatenation of $\mathbf{z}_n$ and $\mathbf{z}_c$. In this work, we use $\mathbf{z}_n \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{d_n})$, where we chose a small value of variance (σ) as 0.10 to make the clusters separated. We use $\mathbf{z}_c$ as a one-hot encoded vector by using the original speaker labels in the training data. Thus, our ClusterGAN training is supervised in nature. The mixture of $\mathbf{z}_n$ and $\mathbf{z}_c$ as the prior enables clustering in the latent space.

c) Inverse mapping network: Inverse mapping from data space to latent space is a non-trivial problem, as it requires the inversion of the generator, which is a multi-layered non-linear model. The work proposed in [51], [52], tackles this issue by solving an optimization problem in $\mathbf{z}$ to recover the latent vectors using $\mathbf{z}^* = \text{argmin}_{\mathbf{z}} L(G(\mathbf{z}), \mathbf{x}) + \lambda\|\mathbf{z}\|_p$, where L is a suitable loss function, λ is a regularization constant and $\|\cdot\|_p$ denotes the norm. However, this optimization is non-convex in $\mathbf{z}$ and there exist multiple $\mathbf{z}$ values to describe a

single real data $\mathbf{x}$ [26], [52]. To mitigate these issues, the stochastic clipping of $\mathbf{z}$ at each iteration step was used in [51]. However, the above approaches are not amenable to clustering. In this work, we train a separate encoder (E) network alongside the GAN network to learn the inverse mapping function of the generator, estimating discriminative latent embeddings for the real data. We fix $\mathbf{z}_c$ and use multiple restarts, each time sampling $\mathbf{z}_n$ from a normal distribution. Moreover, to enforce precise recovery of $\mathbf{z}_n$, we compute the numerical difference between $\mathbf{z}_n$ and corresponding encoder output $\hat{\mathbf{z}}_n$. We empirically found that instead of mean square error, cosine distance is more suitable in the embedding space for distance calculation. The objective function related to this task is

$$\min \text{COS}(G, E) = \frac{1}{m} \sum_{i=1}^{m} \left[1 - \frac{E(G(\mathbf{z}_n^i)) \cdot \mathbf{z}_n^i}{\|E(G(\mathbf{z}_n^i))\| \|\mathbf{z}_n^i\|} \right] \qquad (3)$$

where m is the mini-batch size.

d) Clustering-specific loss: We introduce a clustering-specific loss to learn cluster-friendly representation. For that, we employ cross-entropy (CE) loss, which is computed between $\mathbf{z}_c$ and the soft-max layer output $\hat{\mathbf{z}}_c$ of the encoder network. This loss along with the GAN mini-max objective and the latent variable recovery loss in $\mathbf{z}_n$ encourages clustering in the latent space and also increases discriminative information. We minimize the cross-entropy between the predicted result and the ground truth as

$$\min \text{CE}(G, E) = \frac{1}{m} \sum_{i=1}^{m} \left[p(\mathbf{z}_{c,i}^k) \log p\left(E(G(\mathbf{z}_{c,i}^k)) \right) \right] \qquad (4)$$

where the first term is the empirical probability that the embedding belongs to the k-th speaker, and the second term is the predicted probability that the encoder produced embedding belongs to the k-th speaker.

e) Joint training: The GAN and the encoder networks training in this approach involves joint parameter updates. The final training objective has the following form:

$$\min_{G,E} \max_{D} \left[w_1 \cdot U_{\text{IWGAN}}(D, G) + w_2 \cdot \text{COS}(G, E) + w_3 \cdot \text{CE}(G, E) \right] \qquad (5)$$

Algorithm 1 ClusterGAN training. Default values: $\lambda = 10$, $m = 128$, $n_{\text{critic}} = 5$, $\alpha = 1\text{e}{-4}$, $\beta_1 = 0.5$, $\beta_2 = 0.9$

Require: λ: gradient penalty coefficient; α: learning rate; m: batch size; N_{it}: number of iterations; n_{critic}: number of critic iterations for each generator iteration; α, β_1, β_2: Adam hyper-parameters

1: **for** $it = 1$ to N_{it} **do**
2:      **for** $\tau = 1$ to n_{critic} **do**
3:          Sample $\{\mathbf{x}^{(i)}\}_{i=1}^{m}$, a batch of x-vectors
4:          Update the discriminator parameters by
5:
$$\theta \leftarrow \text{Adam}[\nabla_\theta \{ \frac{1}{m} \sum_{i=1}^{m} w_1 \cdot [D_\theta(\mathbf{x}^{(i)}) - D_\theta(G_\phi(\mathbf{z}^{(i)})) + \lambda \cdot \text{GP}]\}, \theta, \alpha, \beta_1, \beta_2]$$
6:      **end for**
7:      Sample $\{\mathbf{z}^{(i)}\}_{i=1}^{m}$, a batch of latent vectors
8:      Update the generator and encoder parameters by
9:
$$\phi, \psi \leftarrow \text{Adam}[\nabla_{\phi,\psi} \{ \frac{1}{m} \sum_{i=1}^{m} -w_1 \cdot D_\theta(G_\phi(\mathbf{z}^{(i)})) + w_2 \cdot \text{COS}(G_\phi, E_\psi) + w_3 \cdot \text{CE}(G_\phi, E_\psi)\}, \phi, \psi, \alpha, \beta_1, \beta_2]$$
10: **end for**

Weights w_2 and w_3 represent relative significance of preserving continuous and discrete portions of the latent variable. Algorithm 1 lists the whole ClusterGAN training procedure.

2) Meta-learning: Thus far we have discussed the training procedure of ClusterGAN, which is considered as a pre-training part of MCGAN training. In the second phase of MCGAN training, we fine-tune the pre-trained encoder with meta-learning based prototypical loss.

a) Meta-learning using prototypical networks: Prototypical networks, or protonets, apply a simpler inductive bias (in the form of class prototypes) as compared to other metric-learning based methods and achieve state-of-the-art few-shot performance in image classification [27] and natural language processing [53]. The key assumption is that there exists an embedding in which samples from each class cluster around a single prototype representation of that class. Protonets learn a non-linear transformation into an embedding space, where every class is represented by its prototype, sample mean of its support set in the embedding space. During inference, an embedded query sample is assigned to its nearest prototype. The encoder network from ClusterGAN is our protonet.

b) Motivation for fine tuning: The motivation behind fine-tuning the encoder with prototypical loss is that it has good generalization ability at test-time to new classes (unseen during training) given only a handful of examples of each new class [27]. Similar to this setting, in speaker diarization, a trained model for embedding extraction is asked to do clustering among unseen speakers within an audio stream. This is close to a metric learning task, where input audio must be mapped to a discriminative embedding space. Furthermore, a speaker embedding such as x-vector is trained on a speaker classification loss, which is not explicitly designed to optimize embedding similarity. Metric learning related losses such as contrastive loss [54], triplet loss [43] can resolve the above issues. Nonetheless, these methods require careful pair or

triplet selection, which is sometimes time-consuming and performance-sensitive. In this context, prototypical loss trained for learning a metric space to mimic the test scenario might be handy in capturing information related to both generalization and clustering objectives.

c) Episode training: The encoder or the protonet in the MCGAN is trained episodically, where each episode is one mini-batch consisting of N_C categories randomly sampled from total K categories (here, speakers). The mini-batch also contains a labeled set of examples (*support* set S) and unlabeled data (*query* set Q) to predict classes. Consider the support set S of N labeled examples as $S = \{\mathbf{x}_i, y_i\}_{i=1}^{N}$, where each sample $\mathbf{x}_i$ is a D-dimensional x-vector in our case and the corresponding speaker label $y_i \in \{1, \ldots, K\}$. We denote $S_k \subseteq S$ as the set of examples labeled with class k. The protonet learns a non-linear mapping $f_\psi : \mathbb{R}^D \to \mathbb{R}^M$. The M-dimensional prototype of each class is computed as the mean of the embedded support points belonging to that class

$$\mathbf{p}_k = \frac{1}{|S_k|} \sum_{(\mathbf{x}_i, y_i) \in S_k} f_\psi(\mathbf{x}_i) \quad (6)$$

where ψ is the learnable parameters of the encoder.

During training, every query sample $\{\mathbf{x}_j, y_j\} \in Q$ is classified against K speakers based on a soft-max over the distances to each speaker prototypes in the new embedding space:

$$p_\psi(y = y_j | \mathbf{x}_j) = \frac{\exp\big(-d(f_\psi(\mathbf{x}_j), \mathbf{p}_{y_j}) \big)}{\sum_{k'} \exp\big(-d(f_\psi(\mathbf{x}_j), \mathbf{p}_{k'}) \big)} \quad (7)$$

where $d(.)$ represents a distance function. The choice of $d(.)$ can be arbitrary. However, it is shown in [27] that the squared Euclidean distance, which is a particular class of distance function known as Bregman divergence, is good for the clustering problem, and the training algorithm is equivalent to modeling the supports using Gaussian mixture density estimation. Therefore, we also use Euclidean distance as our distance function for proto-learning in the embedding space. The loss function for each mini-batch is the negative log probability for the true class via gradient descent. The prototypical loss within a mini-batch can be written as

$$J_{PTL} = \sum_{\{\mathbf{x}_j, y_j \in Q\}} -\log p(y = y_j | \mathbf{x}_j) \quad (8)$$

d) Extension to whole training set: Suppose the whole training set $\mathcal{D} = \{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_{N_{tr}}, y_{N_{tr}}\}$, where each $y_i \in \{1, \ldots, K\}$. Here, K is the total number of speakers in the training set. We iterate through each episode and in each episode, we randomly sample N_C speakers from total K speakers. For each chosen speaker, N_S number of random samples is selected as the support set and from the rest of the samples of that particular speaker, N_Q number of samples is selected as the query set without replacement. The supports are used to construct the class prototypes using Eq. (6) and the loss is computed with weight updates based on the query samples according to Eq. (7), (8). To increase robustness, instead of using fix total number of speakers, we randomly chose the total number of speakers within an episode. The episodic training procedure is illustrated in Algorithm 2.

Algorithm 2 Meta-learning training for prototypical networks. N_{tr} = number of labeled examples in the training set, K = total number of speakers in the training set, $N_C \leq K$ is the number of speakers per episode, N_S = number of support examples per chosen speaker, N_Q = number of query examples per chosen speaker. $\mathcal{R}(S, N)$ denotes a set of N elements sampled uniformly at random from set S, without replacement.

Require: The whole training set $\mathcal{D} = \bigcup_{k=1}^{K} \mathcal{D}_k$, where $\mathcal{D}_k$ represents the subset of $\mathcal{D}$ containing all elements such that $\{(\mathbf{x}_i, y_i); y_i = k\}$
1: $N_C \leftarrow \mathcal{R}(\{10, 20, \ldots, 150\}, 1)$ ▷ Randomly select total number of speakers in an episode
2: $V \leftarrow \mathcal{R}(\{1, \ldots, K\}, N_C)$ ▷ Randomly select speakers in an episode
3: **for** k in $\{1, \ldots, N_C\}$ **do**
4:     $S_k \leftarrow \mathcal{R}(\mathcal{D}_{V_k}, N_S)$ ▷ Supports
5:     $Q_k \leftarrow \mathcal{R}(\mathcal{D}_{V_k} \backslash S_k, N_Q)$ ▷ Queries
6:     $\mathbf{p}_k \leftarrow \frac{1}{N_S} \sum_{(\mathbf{x}_i, y_i) \in S_k} f_\psi(\mathbf{x}_i)$ ▷ Prototypes
7: **end for**
8: $J_{PTL} \leftarrow 0$
9: **for** k in $\{1, \ldots, N_C\}$ **do**
10:     **for** $(\mathbf{x}_j, y_j)$ in Q_k **do**
11:         $J_{PTL} \leftarrow J_{PTL} + \frac{1}{N_C N_Q}[-d(f_\psi(\mathbf{x}_j), \mathbf{p}_{y_j}) +$
            $\log \sum_{k'} \exp(-d(f_\psi(\mathbf{x}_j), \mathbf{p}_{k'}))]$ ▷ Loss update
12:     **end for**
13: **end for**

D. MCGAN testing

After completion of offline training, only the trained encoder model in MCGAN is used to produce the proposed latent embeddings for the input x-vectors of a given test diarization session. The concatenated latent embeddings ($\mathbf{z}_n$ and $\mathbf{z}_c$) for ClusterGAN or logits for MCGAN are clustered using k-means or NME-SC, and speaker labels of each audio segment are obtained.

E. Normalized Maximum Eigengap Spectral Clustering (NME-SC)

We adopt NME-SC[2] as our spectral clustering method in this paper for speaker diarization evaluation in the unknown number of speakers condition. The NME-SC algorithm can auto-tune parameters of the clustering and also provides improved speaker diarization performance as compared to traditional spectral clustering approaches. As reported in [9], the steps to perform NME-SC are: (a) Construct affinity matrix ($\mathbf{A}$) based on cosine similarity values between the segment embeddings. (b) Binarize $\mathbf{A}$ based on a p-value by converting the p-largest elements in each row of $\mathbf{A}$ to 1 and else to 0. (c) Perform symmetrization on the binarized affinity matrix $\mathbf{A}_p$ to obtain $\bar{\mathbf{A}}_p$ and compute the Laplacian matrix $\mathbf{L}_p$ as $\mathbf{L}_p = \mathbf{D}_p - \bar{\mathbf{A}}_p$, where $\mathbf{D}_p$ is a diagonal matrix and $\mathbf{D}_p = \sum_{j=1}^{n} \bar{\mathbf{A}}_{p,ij}$ (d) Perform eigen decomposition on $\mathbf{L}_p$ and create eigengap vector ($\mathbf{e}_p$). (e) Perform NME analysis to

---

[2]https://github.com/tango4j/Python-Speaker-Diarization

	#Meetings	#Speakers
Train	136	155
Dev	14	17
Eval	12	12

estimate the optimum value $\hat{p}$ and number of clusters k for a given session. (f) Select the k-smallest eigenvalues and the corresponding eigenvectors to construct a matrix $\mathbf{P} \in \mathbb{R}^{n \times k}$. (g) Cluster the row vectors of $\mathbf{P}$ using k-means algorithm. The details of the NME-SC algorithm is precisely described in [9].

IV. DATABASE DESCRIPTION

We evaluate our proposed speaker diarization system on five different diverse databases covering many possible types and domains that are encountered in real-world scenarios.

A. CALLHOME database

CALLHOME contains telephonic conversations with sampling frequency 8 kHz. In speaker diarization literature, NIST 2000 speaker recognition evaluation challenge disk-8 is referred to as CALLHOME [55]. It is a multi-lingual database distributed across six languages: English, Spanish, Arabic, Mandarin, Japanese, and German. The database comprises 500 conversations with the number of speakers in each session varying from 2 to 7. The telephone recordings are of 1 to 10 min duration and their distribution of the number of speakers is given in [4], [5].

B. AMI database

The AMI is a publicly available meeting corpus consists of 171 recordings and 100 hours of data[3]. The meetings are recorded at four different sites (Edinburgh, Idiap, TNO, and Brno). We use the multiple close-talk microphone data post beamforming. For our evaluation, we follow the official speech recognition partition of AMI database with TNO meetings excluded from dev and eval set. The same split is also used in [15], [37]. The train, dev, and eval splits have no speaker overlap and the details are shown in Table I.

C. DIHARD II database

The DIHARD II database is from the DIHARD challenge conducted in 2019. It is a multi-domain database focusing on difficult speaker diarization settings. The database is comprised of diverse recordings collected from domains like meeting speech, restaurant recordings, child language acquisition recordings, YouTube videos, clinical recordings, etc [56]. The DIHARD challenge features two audio input conditions: single-channel and multi-channel. We evaluated our system on single-channel data with reference SAD, which is track 1 in the challenge. Moreover, the database has two subsets: development and evaluation. In this work, speaker diarization performance of proposed and other baseline systems are compared only on the development part of the database. The development set contains 192 recordings typically are of short duration ($<$ 10 min) sessions with the number of speakers in each session varies from 1 to 10.

---

[3]http://groups.inf.ed.ac.uk/ami/download/

D. ADOS and BOSCC databases

The proposed system is also tested on two special child-clinician interaction corpora from the sample of children with autism spectrum disorder (ASD). ADOS (Autism Diagnosis Observation Schedule) is a diagnostic tool based on expert clinical administration and observation that produces a diagnostic algorithm score to inform clinical diagnosis of ASD. ADOS comprises 14 play-based conversational tasks, from within which we select two sub tasks: *Emotion and Social Difficulties and Annoyance* from 272 sessions for our evaluation. Each of these dyadic sessions is of duration < 10 min. BOSCC (Brief Observation of Social Communication Change) is a behavioral observation based autism treatment outcome measure that includes play-based conversational segments of dyadic interaction between a child and an adult (e.g., examiner or caregiver) [57]. In this work, the diarization performance is tested on 24 BOSCC sessions that were collected in a clinical setting. A BOSCC session typically lasts for 12 minutes. The ADOS and BOSCC data considered here are from verbal children and adolescents with autism.

V. EXPERIMENTAL SETUP

A. Speech segmentation

In all the experiments, we have used uniform segmentation (as followed in Kaldi) on the speech intervals specified by the oracle SAD. All the experiments reported in this paper use oracle SAD, which is also a common practice in speaker diarization research [3], [7], [47]. Since our focus is on the effectiveness of proposed embeddings in speaker clustering, we use oracle SAD to eliminate the chance of introducing undesirable error initially due to potential performance uncertainty in automated system SAD. For all the experiments, a sliding window of 1.5 sec duration and overlap of 1 sec is employed to produce speaker-homogeneous segments. Note that in this work no re-segmentation module is applied in the final processing step.

B. x-vector extraction

We use x-vectors from the CALLHOME[4] and Voxceleb[5] recipe as pre-trained audio embeddings for the 8 kHz and 16 kHz data, respectively. The x-vector dimensions are of 128 and 512 for 8 kHz and 16 kHz audio data, respectively.

C. ClusterGAN model specifications

We train two different ClusterGAN models to evaluate diarization performance on various databases. To test speaker diarization performance in CALLHOME which contains 8 kHz telephonic conversations, we train ClusterGAN network in a supervised manner based on AMI train (downsampled to 8 kHz) and switchboard (NIST SRE 2000, disk-6) data. The other model which we employ for diarization performance evaluation on all other databases (AMI, DIHARD II dev, ADOS, BOSCC) containing 16 kHz data is trained on AMI train and ICSI data. We use 60 beamformed ICSI [58] sessions with a total number of 46 speakers. The architectures details of generator (G), discriminator (D) and encoder (E) networks

[4]https://kaldi-asr.org/models/m6
[5]https://kaldi-asr.org/models/m7

TABLE II
CLUSTERGAN ARCHITECTURE DETAILS.

Generator (G)	Discriminator (D)	Encoder (E)
Input: Linear, $\mathbf{z} = (\mathbf{z}_n, \mathbf{z}_c)$ $\in \mathbb{R}^{d_\mathbf{z}}$, $d_n^* = 90$, $d_c^\dagger = 932$ for 8 kHz model and 201 for 16 kHz model	Input: Linear, $\mathbf{x} \in \mathbb{R}^{d_\mathbf{x}}$	Input: Linear, $\hat{\mathbf{x}} \in \mathbb{R}^{d_\mathbf{x}}$
FC‡ 512 ReLU	FC 512 ReLU	FC 512 ReLU
FC 512 ReLU	FC 512 ReLU	FC 512 ReLU
	FC 512 ReLU	FC 1024 ReLU
Output: FC $d_\mathbf{x}$ linear for $\hat{\mathbf{x}}$	Output: FC 1 linear	Output: FC $d_\mathbf{z}$ linear for $\hat{\mathbf{z}}$. Softmax on last d_c to obtain $\hat{\mathbf{z}}_c$

*Dimension of $\mathbf{z}_n$, †Dimension of $\mathbf{z}_c$, ‡Fully-connected

in ClusterGAN are shown in Table II. Moreover, we set the learning rate to 1e-4 and adopt Adam optimization with a mini-batch size of 128 samples to optimize the three networks. We choose the weights w_1, w_2, and w_3 as 1, 10, and 10, respectively, by tuning the diarization error rate (DER) on a held-out set for the 8 kHz model and AMI dev set for the 16 kHz model. It is to be noted that all the above-mentioned model specifications are kept the same for all the experiments reported in this paper.

D. MCGAN specifications

We fine-tune the prototypical network, i.e., the pre-trained encoder in ClusterGAN using Euclidean distance based prototypical loss, which is found to be more effective than cosine distance in [27]. We use the same encoder for embedding extraction for both support and query points; while x-vectors from the training data form the support and queries. We fine-tune the pre-trained encoder by freezing its first two hidden layers and then train it with prototypical loss. We develop support and query set from the same training data that are used to train the ClusterGAN. Instead of using the fixed number of classes to construct all the episodes, we randomly choose the number of classes from 10 to 150 with intervals of 10 per training episode and found this approach is slightly more effective. The number of shots to use in the support set is selected by tuning the DER on the AMI dev set. We fix the number of supports and queries to 10 for all the experiments.

E. Baseline systems

We compare our proposed embeddings with different back-end clustering techniques against several baselines and state-of-the-art diarization systems in five different databases. Since our proposed system incorporates x-vectors as input features, we use Kaldi-based x-vectors with PLDA scoring and AHC clustering as our main baseline system. Furthermore, we show results for x-vector embedding and k-means or spectral clustering (SC) as back-ends, and these are other baseline systems. For a fair comparison, we also report the results of our proposed embeddings with k-means and SC back-ends. We also perform embedding fusion with x-vectors with k-means and SC back-end clustering. Note that for the oracle number of speakers we used fixed tuned p-value binarized SC [9], whereas for the estimated number of speakers we adopt NME-SC [9] for all the experiments in this paper.

F. Performance metrics

We evaluate the proposed speaker diarization system with NIST diarization error rate (DER) [59]. Following the approach described in [59], we use a collar of 0.25 sec for all
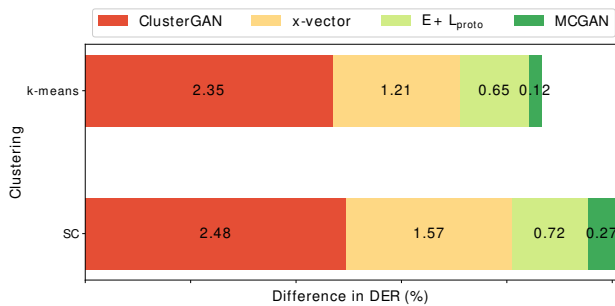
Fig. 3. Difference in DER (%) when the final proposed system (fused) is trained using some components of the proposed system.

TABLE III
RESULTS ON CALLHOME DATABASE FOR MAPD OF PREDICTED SPEAKER NUMBER AND POC IN ESTIMATING SPEAKER NUMBER.

Metric	x-vector	ClusterGAN	MCGAN	Fusion (x-vector + MCGAN)
MAPD	12.54%	11.23%	**9.76%**	10.59%
POC	74.15%	72.14%	**75.55%**	75.35%

the databases DER evaluation, except DIHARD II, where zero collar is used according to the challenge criteria [56]. Since we employ oracle SAD in this work, all the reported DER is exactly speaker confusion and no missed or false alarm speech.

VI. EXPERIMENTAL RESULTS

A. CALLHOME database

1) Ablation study: We perform an ablation study to examine the contribution of each component of our proposed system. To do so, we train a single encoder network with random initialization based on prototypical loss only ($E + L_{proto}$) using x-vectors of the training data as input. We compute DER for different embeddings like ClusterGAN, $E + L_{proto}$, MCGAN, x-vector, x-vector + MCGAN with both k-means and spectral clustering (SC) as back-end clustering. Fig. 3 shows the difference in DER values between each embedding and our final proposed embedding (x-vector + MCGAN) with two mentioned clustering techniques in CALLHOME. The mean difference for all the sessions is shown between each scenario and the final proposed setting.

It is observed from the figure that all the sub-components contribute to improving DER performance. The effect of the components on diarization performance in the CALL-HOME dataset with increasing order is ClusterGAN, x-vector, $E + L_{proto}$, and MCGAN for both the k-means and SC back-ends. For both k-means and SC back-ends, $E + L_{proto}$ is more effective than x-vector. Moreover, the figure shows fine tuning ClusterGAN with prototypical loss (MCGAN) is important in achieving improved DER. It also demonstrates that meta training and embedding fusion are the key components to finally obtain the best results in CALLHOME.

2) Number of predicted speakers: In addition to DER, the mean absolute percentage deviation (MAPD) of the predicted number of speakers and percentage of the correct number of speaker estimation (POC) across all the sessions are also useful metrics in the context of the number of speakers estimation in speaker diarization. The lower the MAPD and the higher the POC is the better speaker estimation. The results on CALLHOME are summarized in Table III. It is evident from the table that MCGAN embeddings are more robust and

TABLE IV
RESULTS ON CALLHOME DATABASE FOR THE BASELINE AND PROPOSED SYSTEMS.

Embedding	Back-end	Avg. DER (%) (oracle SAD, known #speakers)	Avg. DER (%) (oracle SAD, est. #speakers)
x-vector		9.00	8.69
ClusterGAN		10.24	9.83
$E + L_{proto}$	k-means	9.09	8.13
MCGAN		8.72	7.60
x-vector + ClusterGAN		8.98	8.77
x-vector + MCGAN		8.40	7.48
x-vector		6.23	8.32
ClusterGAN		7.62	9.24
$E + L_{proto}$	SC	6.34	7.48
MCGAN		6.01	7.03
x-vector + ClusterGAN		6.22	7.70
x-vector + MCGAN		**5.73**	**6.76**
Wang et al. [6] d-vector	SC	–	12.00
Romero et al. [3] x-vector	PLDA+AHC+VB*	–	9.90
Kaldi x-vector	PLDA+AHC+CV†	7.12	8.39
Zhang et al. [7] d-vector (5-fold)	UIS-RNN+CV	–	7.60
Park et al. [9] Kaldi x-vector	NME-SC	–	7.29

*Variational Bayes re-segmentation, †Cross-validation

accurate in estimating the number of speakers than Cluster-GAN embeddings and x-vectors. The performance of fused embeddings is slightly worse than MCGAN. Therefore, it is expected that MCGAN will perform better than ClusterGAN and x-vector in the estimated number of speakers condition.

3) Overall performance evaluation: In this section, we present the experimental results on the whole CALLHOME evaluation set by using the tuned parameters of the different versions of our proposed diarization system. We compare the proposed system with other baselines and recent state-of-the-art diarization methods. The experimental results for both known and unknown numbers of speakers are reported in Table IV. Note that for known or oracle number of speakers we use fix p-value which is tuned on Kaldi CALLHOME-1 held-out set and apply it to CALLHOME-2 and vice versa.

From Table IV column 3, we observe that for known number of speakers, the ClusterGAN embedding does not outperform x-vectors for both k-means and SC. However, we see that MCGAN embeddings which are extracted after fine-tuning the pre-trained encoder with prototypical loss provide superior performance over x-vectors for both k-means and SC back-ends. MCGAN reduces average DER of ClusterGAN from 10.24% to 8.72% and from 7.62% to 6.01% for k-means and SC, respectively. Therefore, fine-tuning the protonet (E) with meta-learning related prototypical loss is useful for better generalization. To corroborate this fact, the average DER of an encoder initialized with random weights and trained with prototypical loss ($E + L_{proto}$) on the training data is also calculated and the corresponding results are given in Table IV row 3 and 9. It is observed from the table that proposed MCGAN embeddings reduce DER of $E + L_{proto}$ from 9.09% to 8.72% and from 6.34% to 6.01% respectively for k-means and SC. Finally, we obtain further improvement in DER by incorporating embedding fusion between x-vector and MCGAN embeddings. We achieve the best DER of 5.73% for the known number of speakers and SC back-end, which is significantly better than the Kaldi x-vector state-of-the-art (average DER 7.12%) and also superior to the x-vector with SC (average DER 6.23%). The relative improvement of our final proposed system over Kaldi state-of-the-art is 19.52%
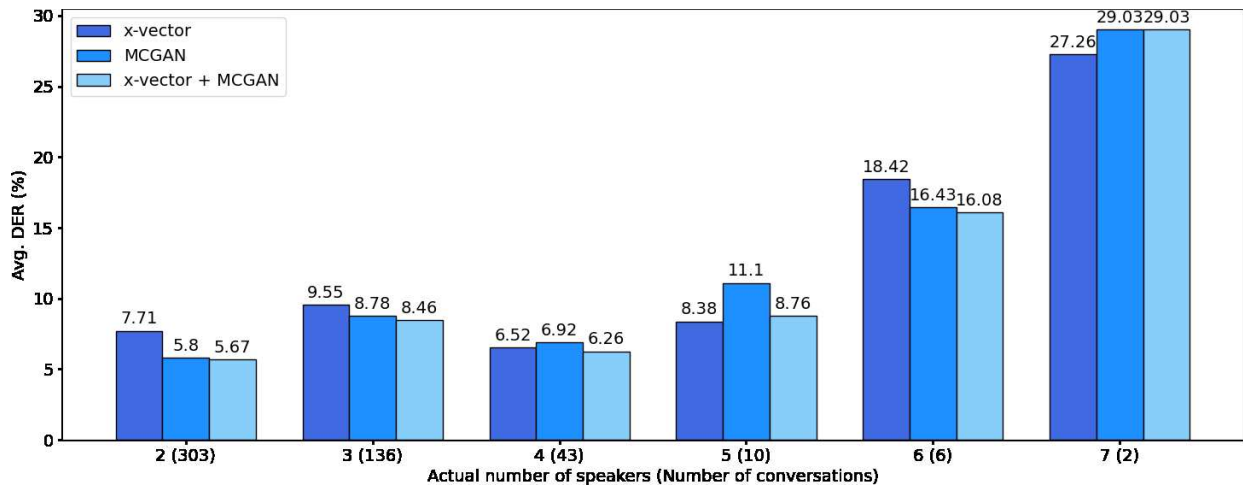
Fig. 4. Avg. DER (%) analysis with respect to number of speakers in a diarization session.

for known number of speakers.

We show the diarization performance of all the systems for the estimated number of speakers in Table IV column 4. The number of speakers for k-means and SC is estimated using NME-SC. From Table IV column 4, we see a similar trend in performance for estimated number of speakers. The biggest improvement in DER for proposed embeddings comes from MCGAN, embedding fusion, and most importantly the SC. For the same back-end setting, the proposed MCGAN and fused embeddings are better than x-vector. It is important to note that surprisingly in many of the settings (except SC) we obtain reduced DER for the automatically estimated number of speakers case than for oracle number of speakers. This could be attributed to the fact that even though the number of clusters may be correct for the oracle case there might be inherent speaker confusions, whereas, for the estimated number of speakers, the clusters based on data-driven estimation may be purer even if the estimated number of clusters is not exactly correct. The embedding fusion between x-vector and MCGAN with SC back-end yields the best DER value of 6.76% for the estimated number of speakers with a relative improvement of 19.43% over the Kaldi x-vector system. The next best system–MCGAN with SC–produces a DER of 7.03%, which is also significantly better than the Kaldi x-vector and x-vector with SC back-end. We also present the recent best system's results that are reported in the literature on the CALLHOME evaluation set. Many of these systems use cross-validation to train or adapt their systems. However, without using any cross-validation, the proposed system outperforms all the recent diarization systems in CALLHOME.

4) Analysis of Experimental Results: We first break down the average DER on CALLHOME database according to the number of speakers. The corresponding DERs are plotted in a group bar plot in Fig. 4 for x-vector, MCGAN and fused (x-vector + MCGAN) embeddings with NME-SC back-end and estimated number of speakers. It is evident from the figure that our proposed MCGAN and fused embeddings achieve significantly better DER values than x-vector for two and three speaker cases. For four and five speakers, x-vector is better than MCGAN. However, the fused embeddings provide better
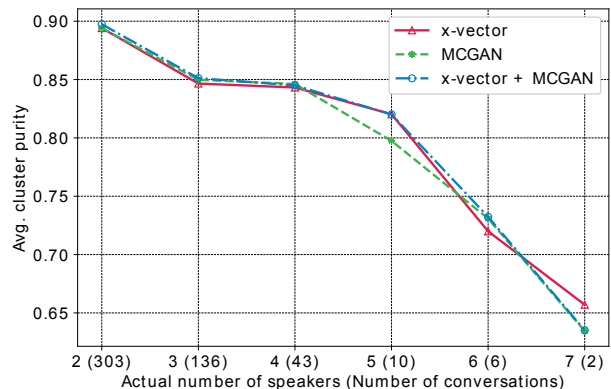

Fig. 5. Average cluster purity analysis with respect to number of speakers in a diarization session.

performance than x-vectors for most of the speaker conditions (two, three, four, and six) and this covers majority of the conversations in the dataset. In the seven-speaker condition, the fused system is not able to outperform x-vector, which is possibly an anomaly since the number of sessions containing seven speakers is only two in this database. The reason behind obtaining better results using fused embeddings with SC is attributed to the complementary merits of the x-vector and MCGAN embeddings, and the modeling power of the NME-SC algorithm on embeddings.

In light of this, it would be reasonable to consider another performance metric related to the clustering mechanism. Fig. 5 shows the average cluster purity with respect to changes in the number of speakers in CALLHOME for x-vector, MCGAN, and fused (x-vector + MCGAN) embeddings with NME-SC back-end and estimated number of speakers. We can see that the fused embeddings provide better cluster purity than x-vector and MCGAN for all the number of speaker cases except 7, which has only two sessions. We observe noticeable improvement for MCGAN and fused over x-vector for the six-speaker case and little improvement in two and three speaker conditions, a result consistent with our DER analysis with respect to the number of speakers.

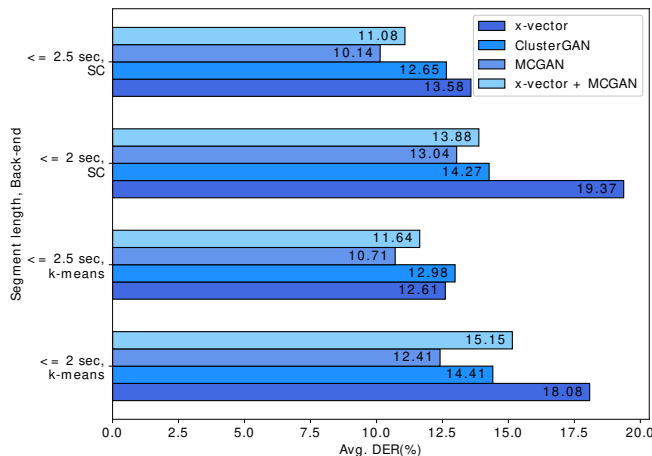We extend the analysis by checking the effectiveness of

Fig. 6. Average DER (%) analysis on short speech segment diarization sessions.

TABLE V
RESULTS ON AMI DEV AND EVAL SETS FOR THE BASELINE AND PROPOSED SYSTEMS.

Embedding	Back-end	Avg. DER (%) (oracle SAD, known #speakers)		Avg. DER (%) (oracle SAD, est. #speakers)	
		Dev	Eval	Dev	Eval
x-vector	k-means	11.94	11.45	12.64	12.26
ClusterGAN		7.64	7.69	11.34	11.51
MCGAN		5.84	6.14	7.09	6.09
x-vector + ClusterGAN		6.62	6.46	9.57	8.63
x-vector + MCGAN		5.64	5.48	6.47	8.76
x-vector	SC	7.32	6.88	6.42	6.23
ClusterGAN		**3.86**	3.91	6.41	8.16
MCGAN		5.72	4.49	5.10	5.38
x-vector + ClusterGAN		3.93	**3.60**	6.21	**2.87**
x-vector + MCGAN		5.49	4.23	**5.02**	4.92
Kaldi x-vector	PLDA+AHC	11.65	11.34	11.08	10.37

our proposed system in a more challenging practical scenario namely diarization in short speech segment case. Shorter segments usually provide low-quality speaker embeddings. To carry out this analysis, we chose conversations from the CALLHOME evaluation set that have a majority number of short duration ($<=$ 2 sec and $<=$ 2.5 sec) speech segments. Here, we select sessions that have more than 80% of short speech segments in the entire session. We find that number of such sessions is 58 ($<=$ 2 sec) and 129 ($<=$ 2.5 sec), respectively. We compute and plotted the mean DER of the selected sessions in Fig. 6 for x-vector, ClusterGAN, MCGAN, and fused (x-vector + MCGAN) embeddings with k-means and spectral clustering and estimated number of speakers. It is clear from the figure that among the four embeddings, MCGAN embedding produces the lowest average DER for short speech segment sessions compared to x-vector, ClusterGAN, and fused embeddings, and for both the clustering techniques. The fused and ClusterGAN embeddings yield better performance than x-vector for most of the cases. We obtain worse DER values for $<=$ 2 sec segments than $<=$ 2.5 segments, which is not surprising. Finally, we can conclude that MCGAN embedding is more robust than the other embeddings in short speech segment scenarios.

B. AMI database

1) Performance comparison: Herein we evaluate the diarization performance of several proposed systems and compare them with Kaldi state-of-the-art and other alternative baselines on the AMI meetings corpus. We use the model trained on AMI train and ICSI data to generate our proposed speaker embeddings.

Results presented in Table V show speaker diarization performance on both the AMI dev and eval sets for oracle SAD with both known and estimated number of speakers. The hyper-parameters for ClusterGAN and MCGAN models and the p-value for spectral clustering are tuned on the dev set and applied on the dev and eval set. From Table V column 3, we see that for the known number of speakers, all the proposed embeddings (ClusterGAN, MCGAN, x-vector + ClusterGAN, x-vector + MCGAN) significantly outperform x-vector for both k-means and SC back-ends. This is true for both AMI

dev and eval sets. Moreover, for various proposed settings, the embedding fusion reduces the DER over individual embeddings and SC further lowers the DER as compared to the k-means counterpart. The fused x-vector + ClusterGAN with SC achieves the best DER value of 3.60% on AMI eval set for the known number of speakers case. This indicates that the proposed embeddings have complementary information to x-vector embeddings. In addition, the fused embeddings with SC outperform Kaldi x-vector baseline system by a significant margin.

We show results for all the systems on AMI dev and eval sets for the estimated number of speakers in Table V column 4. The number of speakers for both k-means and SC is estimated by using the NME-SC algorithm. For the Kaldi x-vector baseline, the number of speakers in a session is estimated based on thresholding on the PLDA scores [3]. It is clear from Table V column 4 that similar to the oracle case, the proposed embeddings are superior to x-vector in dev and eval set under k-means clustering. MCGAN yields better performance compared to the ClusterGAN. As seen previously, MCGAN embeddings are more robust in the estimated number of speakers condition than ClusterGAN. Moreover, as expected, the fused embeddings further improve diarization performance for the k-means back-end. On the other hand, spectral clustering boosts diarization performance further for all the embeddings. This ensures the effectiveness of NME-SC over k-means. Besides, the embedding fusion provides further reduction in DER for SC back-end. Finally, using our proposed fused system, we achieve significantly better performance on AMI dev and eval set with absolute DER of 5.02% and 2.87%, respectively, outperforming Kaldi x-vector baseline diarization system.

2) t-SNE visualization: Fig. 7 shows the t-SNE visualization of x-vector, ClusterGAN, MCGAN, x-vector + ClusterGAN and x-vector + MCGAN embeddings. While the proposed ClusterGAN and MCGAN embeddings exhibit compact clusters than x-vectors, MCGAN shows better separation among the clusters than ClusterGAN. It is also evident from the figure that fused embeddings exhibit compact and well-separated clusters. Moreover, in this particular example, x-vector + MCGAN shows near-perfect clustering among the four classes, which is slightly better than x-vector + ClusterGAN.
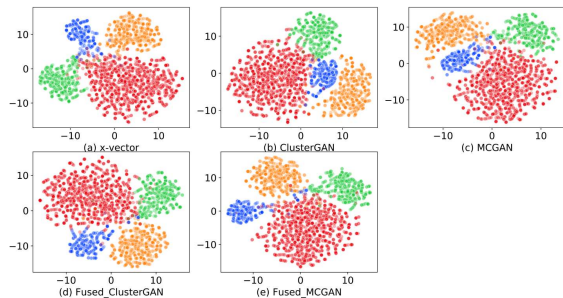
Fig. 7. t-SNE visualization of (a) x-vector, (b) ClusterGAN, (c) MCGAN (d) x-vector + ClusterGAN, and (e) x-vector + MCGAN embeddings of IS1008a AMI session. This AMI session contains four speakers and each speaker is represented by different colours in the figure.

C. DIHARD II database

1) Performance comparisons: To investigate the effectiveness of the proposed embeddings for speaker diarization in challenging multi-domain settings, we evaluated and report the average DER values on the DIHARD II development database in Table VI. We also compare different versions of proposed embeddings against x-vector with different clustering techniques and existing state-of-the-art systems. It is to be noted that the proposed embeddings are based on the same model that was used for AMI database evaluation. This is to check the robustness of our embeddings in real-world noisy scenarios without training explicitly using separate noisy data or data augmentation. Note that for the known number of speakers since we used fixed p-value for SC in this paper, we tune it on the same dev set for all the embeddings. It is seen from Table VI column 3 that x-vector produces the best DER value as compared to ClusterGAN and MCGAN for k-means clustering. This is reasonable since the x-vector pre-trained model is based on Voxceleb data with data augmentation. However, embedding fusion with x-vectors improves the performance further for both the proposed embeddings. In contrast to k-means clustering, spectral clustering provides a larger boost in performance for all the embeddings. We attain lower DER value for ClusterGAN than x-vector and MCGAN for known number of speakers and SC back-end.

We report the DER values for the estimated number of speakers using NME-SC in Table VI column 4. While the performance is not good for k-means with oracle number of speakers, we obtain surprisingly better DER values for the estimated number of speakers. It is seen from the table that MCGAN is superior to ClusterGAN in the estimated number of speakers scenario. However, individually x-vector is better than ClusterGAN and MCGAN in this database. Nonetheless, both the fused embeddings outperform x-vector for k-means back-end. With SC back-end, we achieve significant improvement in performance for all the embeddings. We attain the best DER value of 17.75% by using x-vector + ClusterGAN embedding and SC back-end. For comparison, we also report the performance of several existing DIHARD challenge II submissions. The challenge top system by BUT achieves a DER value of 18.09% on the DIHARD II dev set [60]. However, it is mentioned in the paper that in their system, PLDA was adapted on the same development

TABLE VI
RESULTS ON DIHARD II DEVELOPMENT SET FOR THE BASELINE AND PROPOSED SYSTEMS.

Embedding	Back-end	Avg. DER (%) (oracle SAD, known #speakers)	Avg. DER (%) (oracle SAD, est. #speakers)
x-vector		27.12	23.38
ClusterGAN		29.68	26.03
MCGAN	k-means	28.98	25.71
x-vector + ClusterGAN		27.76	22.87
x-vector + MCGAN		26.59	22.53
x-vector		19.86	18.88
ClusterGAN		19.08	21.84
MCGAN	SC	19.87	21.16
x-vector + ClusterGAN		**17.69**	**17.75**
x-vector + MCGAN		18.86	18.59
Kaldi x-vector	PLDA+AHC	–	25.59
Kaldi x-vector	PLDA(adapted)+AHC	–	23.82
[17] Kaldi x-vector(fusion)	SC	–	21.82
[61] Resnet embed.	SC+VB*+OD†	–	18.84
[60] x-vector	PLDA+AHC+VB+OD	–	18.09

*Variational Bayes re-segmentation, †Overlap detection

set. Therefore, it is worthwhile to note that without using PLDA and several additional processing steps like speech enhancement, overlap detection, and variational Bayes (VB) re-segmentation our proposed system outperforms the top system in the DIHARD II development set. Furthermore, the proposed system is significantly better than the Kaldi x-vector diarization system, which was the baseline in the challenge. Thus, the proposed embeddings although extracted from the model trained on AMI train and ICSI data, are promising in terms of generalization, have complementary information to x-vectors, and can yield state-of-the-art performance on a challenging multi-domain database in an embedding fusion set up with the spectral clustering back-end.

2) DER analysis according to the domains: To understand how our proposed embeddings with spectral clustering behave in each specific domain of DIHARD II dev set, we split the DER according to the context of the database. The results shown in Table VII indicate high variability in performance across the domains. The proposed embeddings (ClusterGAN and MCGAN) individually are not able to outperform x-vectors except on audiobooks data, which contains only one speaker. However, the fused embeddings offer promising performance on most of the domains compared to the x-vectors. The worst performing domains for our embeddings are restaurant, webvideo, and child. The metadata analysis of DIHARD II dev set in [61] shows that restaurant sessions are highly noisy and also contain a large number of speakers. In addition, restaurant and webvideo data comprise significant speech overlap. On the other hand, although the child data contains less amount of overlap, the observed worse performance is because the children in the sessions are 6-18 months old, and have high variability in their speech; moreover, more than two speakers are present in those sessions. We also compare our results with a recently proposed diarization system on DIHARD II data, which was also the second best performing system in the DIHARD II challenge [61]. It is intriguing to note that our embeddings perform well in the meeting domain despite the presence of high overlapping speech error. This is possible because the proposed embeddings were trained on meeting data. However, x-vectors that are extracted from Voxceleb trained model also perform well. The other domains

TABLE VII
DIARIZATION PERFORMANCE OF THE PROPOSED AND BASELINE SYSTEMS IN EACH SPECIFIC DOMAIN ON THE DIHARD II DEV SET.

System	audiobooks *(12)	broadcast _interview (12)	maptask (23)	socio_lab (16)	socio_field (12)	court (12)	clinical (24)	child (23)	webvideo (32)	meeting (14)	restaurant (12)
x-vector + NME-SC	5.92	**4.56**	9.23	**7.87**	12.17	5.60	20.20	31.65	**33.74**	10.20	42.68
ClusterGAN + NME-SC	1.98	4.75	9.50	8.75	15.54	4.67	24.48	31.77	37.37	24.94	54.04
MCGAN + NME-SC	1.59	4.86	10.14	11.99	11.64	**4.66**	29.33	34.26	36.38	12.50	44.41
(x-vector + ClusterGAN) + NME-SC	2.29	4.68	9.10	8.20	**9.20**	5.31	**16.37**	**29.69**	34.66	**9.90**	**40.46**
(x-vector + MCGAN) + NME-SC	3.64	4.77	**8.92**	8.63	12.28	4.78	18.46	31.06	34.42	11.24	42.10
Lin et al. [61]	**0.0**	5.52	11.32	12.08	14.46	12.12	17.42	37.38	34.90	25.73	43.00

* The number of sessions within each domain

where we obtain noticeable improvements over the existing system include court, child, socio_field, and socio_lab.

D. ADOS and BOSCC databases

Finally, we evaluate the proposed method on two child-clinician interaction corpora from the domain of Autism Spectrum Disorder: ADOS and BOSCC. The diarization results are presented in Table VIII. We observe from the table that the Kaldi x-vector diarization system (last row) does not perform well on these two databases. The most probable reason behind this is that the PLDA model is trained on Voxceleb data and thus creating a significant domain mismatch. However, the x-vectors with k-means and SC perform reasonably well on both ADOS and BOSCC data than the Kaldi x-vector system. Among the proposed embeddings, ClusterGAN is superior to MCGAN both individually and also when fused with x-vectors. This is attributed to the better performance of ClusterGAN over MCGAN in the known number of speakers condition in general. A significant reduction in DER is seen while SC is employed as the clustering mechanism. The best achieved DER on ADOS and BOSCC datasets is 6.74% and 9.26%, respectively, and this is for the x-vector + ClusterGAN with SC system. We obtain a relative improvement of 53.06% and 57.31% over Kaldi x-vector on the ADOS and BOSCC databases respectively. We note that although we expect better generalization from MCGAN due to meta-learning, Cluster-GAN emerges as useful in these known number of (dyadic) speaker conditions, i.e., child and adult interlocutors.

VII. CONCLUSIONS

We proposed new speaker embeddings by exploiting the latent space of GANs using ClusterGAN and by making the encoder in the ClusterGAN more robust and generalizable with the help of prototypical loss fine-tuning. We benchmarked the proposed embeddings individually and also fused with x-vectors within the speaker diarization framework. We investigated the effectiveness of the proposed embeddings by extensively evaluating them for speaker diarization across five different databases. We obtain a relative improvement of 19.43%, 72.32%, 30.64%, 53.06%, and 57.31% over the Kaldi x-vector baseline on CALLHOME, AMI-eval, DIHARD II dev, ADOS, and BOSCC databases respectively. The key findings of this work can be summarized as follows:

TABLE VIII
RESULTS ON ADOS AND BOSCC DATABASES FOR THE BASELINE AND PROPOSED SYSTEMS.

Embedding	Back-end	Avg. DER (%) (oracle SAD) on ADOS	Avg. DER (%) (oracle SAD) on BOSCC
x-vector		12.35	14.73
ClusterGAN		10.21	10.59
MCGAN	k-means	9.71	14.67
x-vector + ClusterGAN		8.70	10.52
x-vector + MCGAN		9.10	13.22
x-vector		8.51	11.99
ClusterGAN		6.75	9.32
MCGAN	SC	9.96	13.21
x-vector + ClusterGAN		**6.74**	**9.26**
x-vector + MCGAN		9.18	12.17
Kaldi x-vector	PLDA+AHC	14.36	21.69

- MCGAN embeddings outperform x-vectors and ClusterGAN embeddings significantly on telephonic data for both known and automatically estimated number of speaker conditions with both k-means and SC back-end. They also perform better than ClusterGAN in the estimated number of speaker condition on meeting and multi-domain datasets.
- Analysis suggests that MCGAN embeddings are robust in the number of speakers estimation and in diarizing sessions which have significant presence of short speech segments when compared to x-vectors, ClusterGAN and fused embeddings.
- Embedding fusion of x-vectors and the proposed embeddings improves diarization performance consistently for all the corpora considered. Therefore, we speculate that both the proposed embeddings have complementary information to the x-vectors. The proposed fused embeddings with NME-SC outperform the Kaldi x-vector system and emerges as the top-performing system on the challenging multi-domain DIHARD II dev set.

In the future, it would be worthwhile to investigate speech spectrograms directly instead of pre-trained embeddings as the input. The usage of other existing meta-learning algorithms will also be explored in the context of speaker diarization.

REFERENCES

[1] X. Anguera *et al.*, "Speaker diarization: A review of recent research," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 20, no. 2, pp. 356–370, 2012.

[2] D. Vijayasenan, F. Valente, and H. Bourlard, "An information theoretic approach to speaker diarization of meeting data," *IEEE Trans. Audio, Speech, and Lang. Process.*, vol. 17, no. 7, pp. 1382–1393, 2009.

[3] D. Garcia-Romero *et al.*, "Speaker diarization using deep neural network embeddings," in *Proc. ICASSP*, 2017, pp. 4930–4934.

[4] S. H. Shum, N. Dehak, R. Dehak, and J. R. Glass, "Unsupervised methods for speaker diarization: An integrated and iterative approach," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 21, no. 10, pp. 2015–2028, 2013.

[5] M. Senoussaoui, P. Kenny, T. Stafylakis, and P. Dumouchel, "A study of the cosine distance-based mean shift for telephone speech diarization," *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 22, no. 1, pp. 217–227, 2014.

[6] Q. Wang *et al.*, "Speaker diarization with LSTM," in *Proc. ICASSP*, 2018, pp. 5239–5243.

[7] A. Zhang *et al.*, "Fully supervised speaker diarization," in *Proc. ICASSP*, 2019, pp. 6301–6305.

[8] G. Sell *et al.*, "Diarization is hard: Some experiences and lessons learned for the JHU team in the inaugural DIHARD challenge," in *Proc. Interspeech*, 2018, pp. 2808–2812.

[9] T. J. Park, K. J. Han, M. Kumar, and S. Narayanan, "Auto-tuning spectral clustering for speaker diarization using normalized maximum eigengap," *IEEE Sign. Process. Lett.*, 2019.

[10] Q. Lin *et al.*, "LSTM based similarity measurement with spectral clustering for speaker diarization," *arXiv preprint arXiv:1907.10393*, 2019.

[11] N. Dehak *et al.*, "Front-end factor analysis for speaker verification," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 4, pp. 788–798, 2010.

[12] L. Wan *et al.*, "Generalized end-to-end loss for speaker verification," in *Proc. ICASSP*, 2018, pp. 4879–4883.

[13] Z. Zajíc, M. Hrúz, and L. Müller, "Speaker diarization using convolutional neural network for statistics accumulation refinement," in *Proc. Interspeech*, 2017, pp. 3562–3566.

[14] D. Snyder *et al.*, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. ICASSP*, 2018, pp. 5329–5333.

[15] G. Sun, C. Zhang, and P. C. Woodland, "Speaker diarisation using 2D self-attentive combination of embeddings," in *Proc. ICASSP*, 2019, pp. 5801–5805.

[16] H. Ning *et al.*, "A spectral clustering approach to speaker diarization," in *Proc. ICSLP*, 2006.

[17] T. J. Park *et al.*, "The Second DIHARD challenge: System Description for USC-SAIL Team," in *Proc. Interspeech*, 2019, pp. 998–1002.

[18] D. Dimitriadis, "Enhancements for audio-only diarization systems," *arXiv preprint arXiv:1909.00082*, 2019.

[19] P. A. Mansfield, et al., "Links: A high-dimensional online clustering method," *arXiv preprint arXiv:1801.10123*, 2018.

[20] K. J. Han, S. Kim, and S. S. Narayanan, "Strategies to improve the robustness of agglomerative hierarchical clustering under data source variation for speaker diarization," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 16, no. 8, pp. 1590–1601, 2008.

[21] E. Fini and A. Brutti, "Supervised online diarization with sample mean loss for multi-domain data," in *Proc. ICASSP*, 2020, pp. 7134–7138.

[22] Y. Fujita *et al.*, "End-to-End Neural Speaker Diarization with Permutation-Free Objectives," in *Proc. Interspeech*, 2019, pp. 4300–4304.

[23] ——, "End-to-end neural speaker diarization with self-attention," *arXiv preprint arXiv:1909.06247*, 2019.

[24] J. Xie, R. Girshick, and A. Farhadi, "Unsupervised deep embedding for clustering analysis," in *Proc. ICML*, 2016, pp. 478–487.

[25] M. Pal *et al.*, "Speaker diarization using latent space clustering in generative adversarial network," in *Proc. ICASSP*, 2020, pp. 6504–6508.

[26] S. Mukherjee, H. Asnani, E. Lin, and S. Kannan, "ClusterGAN: Latent space clustering in generative adversarial networks," in *Proc. AAAI*, vol. 33, 2019, pp. 4610–4617.

[27] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proc. NIPS*, 2017, pp. 4077–4087.

[28] J. Wang *et al.*, "Centroid-based deep metric learning for speaker recognition," in *Proc. ICASSP*. IEEE, 2019, pp. 3652–3656.

[29] E. Aljalbout *et al.*, "Clustering with deep learning: Taxonomy and new methods," *arXiv preprint arXiv:1801.07648*, 2018.

[30] Z. Jiang *et al.*, "Variational deep embedding: An unsupervised and generative approach to clustering," *arXiv preprint arXiv:1611.05148*, 2016.

[31] X. Chen *et al.*, "InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets," in *Proc. NIPS*, 2016, pp. 2172–2180.

[32] Y. Yu and W.-J. Zhou, "Mixture of GANs for Clustering." in *Proc. IJCAI*, 2018, pp. 3047–3053.

[33] B. Yang *et al.*, "Towards k-means-friendly spaces: Simultaneous deep learning and clustering," in *Proc. ICML*. JMLR. org, 2017, pp. 3861–3870.

[34] J. Yang, D. Parikh, and D. Batra, "Joint unsupervised learning of deep representations and image clusters," in *Proc. CVPR*, 2016, pp. 5147–5156.

[35] U. Shaham *et al.*, "Spectralnet: Spectral clustering using deep neural networks," *arXiv preprint arXiv:1801.01587*, 2018.

[36] K. Ghasedi *et al.*, "Balanced self-paced learning for generative adversarial clustering network," in *Proc. CVPR*, 2019, pp. 4391–4400.

[37] Q. Li *et al.*, "Discriminative neural clustering for speaker diarisation," *arXiv preprint arXiv:1910.09703*, 2019.

[38] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," 2016.

[39] K. Chen and A. Salman, "Learning speaker-specific characteristics with a deep neural architecture," *IEEE Trans. on Neural Networks*, vol. 22, no. 11, pp. 1744–1756, 2011.

[40] A. Santoro *et al.*, "Meta-learning with memory-augmented neural networks," in *Proc. ICML*, 2016, pp. 1842–1850.

[41] O. Vinyals *et al.*, "Matching networks for one shot learning," in *Proc. NIPS*, 2016, pp. 3630–3638.

[42] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition," in *ICML deep learning workshop*, vol. 2. Lille, 2015.

[43] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proc. CVPR*, 2015, pp. 815–823.

[44] J. S. Chung *et al.*, "In defence of metric learning for speaker recognition," *arXiv preprint arXiv:2003.11982*, 2020.

[45] H. Bredin, "Tristounet: triplet loss for speaker turn embedding," in *Proc. ICASSP*, 2017, pp. 5430–5434.

[46] N. R. Koluguri *et al.*, "Meta-learning for robust child-adult classification from speech," in *Proc. ICASSP*, 2020, pp. 8094–8098.

[47] M. Diez *et al.*, "Analysis of speaker diarization based on Bayesian HMM with eigenvoice priors," *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 28, pp. 355–368, 2019.

[48] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. NIPS*, 2014, pp. 2672–2680.

[49] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," *arXiv preprint arXiv:1701.07875*, 2017.

[50] I. Gulrajani *et al.*, "Improved training of Wasserstein GANs," in *Proc. NIPS*, 2017, pp. 5767–5777.

[51] Z. C. Lipton and S. Tripathi, "Precise recovery of latent vectors from generative adversarial networks," *arXiv preprint arXiv:1702.04782*, 2017.

[52] A. Creswell and A. A. Bharath, "Inverting the generator of a generative adversarial network," *IEEE Trans. on neural networks and learning systems*, 2018.

[53] M. Yu *et al.*, "Diverse few-shot text classification with multiple metrics," *arXiv preprint arXiv:1805.07513*, 2018.

[54] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *Proc. CVPR*, vol. 1. IEEE, 2005, pp. 539–546.

[55] A. F. Martin and M. A. Przybocki, "Speaker recognition in a multi-speaker environment," in *Proc. Speech Commun. and Tech.*, 2001.

[56] N. Ryant *et al.*, "Second DIHARD challenge evaluation plan," *Linguistic Data Consortium, Tech. Rep*, 2019.

[57] R. Grzadzinski *et al.*, "Measuring changes in social communication behaviors: preliminary development of the brief observation of social communication change (BOSCC)," *Journal of autism and developmental disorders*, vol. 46, no. 7, pp. 2464–2479, 2016.

[58] A. Janin *et al.*, "The ICSI meeting corpus," in *Proc. ICASSP*, vol. 1, 2003, pp. I–I.

[59] J. G. Fiscus *et al.*, "The Rich Transcription 2006 spring meeting recognition evaluation," in *International Workshop on Machine Learning for Multimodal Interaction*. Springer, 2006, pp. 309–322.

[60] F. Landini *et al.*, "BUT System Description for DIHARD Speech Diarization Challenge 2019," *arXiv preprint arXiv:1910.08847*, 2019.

[61] Q. Lin *et al.*, "DIHARD II is Still Hard: Experimental Results and Discussions from the DKU-LENOVO Team," *arXiv preprint arXiv:2002.12761*, 2020.