



# Cosine-Distance Virtual Adversarial Training for Semi-Supervised Speaker-Discriminative Acoustic Embeddings

Florian L. Kreyssig & Philip C. Woodland

Cambridge University Engineering Dept., Trumpington St., Cambridge, CB2 1PZ U.K.

{flk24, pcw}@eng.cam.ac.uk

## Abstract

In this paper, we propose a semi-supervised learning (SSL) technique for training deep neural networks (DNNs) to generate speaker-discriminative acoustic embeddings (speaker embeddings). Obtaining large amounts of speaker recognition training data can be difficult for desired target domains, especially under privacy constraints. The proposed technique reduces requirements for labelled data by leveraging unlabelled data. The technique is a variant of virtual adversarial training (VAT) [1] in the form of a loss that is defined as the robustness of the speaker embedding against input perturbations, as measured by the cosine-distance. Thus, we term the technique cosine-distance virtual adversarial training (CD-VAT). In comparison to many existing SSL techniques, the unlabelled data does not have to come from the same set of classes (here speakers) as the labelled data. The effectiveness of CD-VAT is shown on the 2750+ hour VoxCeleb data set, where on a speaker verification task it achieves a reduction in equal error rate (EER) of 11.1% relative to a purely supervised baseline. This is 32.5% of the improvement that would be achieved from supervised training if the speaker labels for the unlabelled data were available.

**Index Terms:** semi-supervised, speaker embeddings, d-vector, speaker verification

## 1. Introduction

Speaker-discriminative acoustic embeddings (or just speaker embeddings) derived through deep learning techniques have become the state-of-the-art for learning speaker representations [2, 3] to be used for tasks such as speaker recognition, speaker verification or speaker diarisation [2, 4, 5, 6]. Previously i-vectors [7] based on factor analysis were widely used.

The neural networks used to generate speaker embeddings are typically trained on a speaker classification task, for which the input is the acoustic feature sequence of an utterance and the output is the speaker label of that utterance [2, 8]. By taking the output of a layer of this neural network (often the penultimate layer) as an embedding, a fixed dimensional vector can be generated for any given input utterance. This vector is speaker-discriminative due to the training objective. It has been found that such speaker embeddings can be used to discriminate between speakers that are not present in the training data.

The quality of these embeddings will improve with the amount of training data and with the number of speakers in the training data, assuming the data comes from the target domain. However, the acquisition of enough suitable speaker classification data for the exact conditions one desires can be difficult. This is especially true as the regulations around identifiable user data tighten<sup>1</sup>. Under these constraints it is useful to use audio

data with associated speaker labels together with de-identified (unlabelled) data to train speaker embedding generators.

This paper proposes a method that enables semi-supervised learning (SSL) of speaker embeddings. In comparison to many SSL methods in machine learning, the proposed method does not assume that the labelled and unlabelled data comes from the same classes (here speakers). Therefore, a small amount of labelled data from a small number of speakers can be complemented by a larger amount of data from a large number of speakers. The proposed method is a newly derived sibling to virtual adversarial training (VAT) [1] which is an SSL method for classification tasks. Vanilla VAT assumes the labelled and unlabelled data to come from the same set of classes. This paper, however, attempts to utilise unlabelled data that comes from a completely different set of classes (here speakers).

The proposed SSL technique, termed cosine-distance virtual adversarial training (CD-VAT), works by adding an additional loss to the standard supervised training loss. The loss is defined as the cosine-distance between a speaker embedding generated for an utterance and the embedding generated for the same utterance, which was perturbed by an *adversarial* noise that maximally increases the cosine-distance to the original, unperturbed, embedding. The loss is computed for every data point in the labelled and unlabelled data sets, thus smoothing the embedding generator with respect to (w.r.t) the input for all data points lying on the data manifold. It can, therefore, be seen as a regularisation technique that is informed by the unlabelled data, which constrains the neural network to learn embeddings that generalise well to unseen speakers.

This paper is organised as follows. Section 2 describes the CD-VAT loss and how the adversarial noise is computed. In Sec. 3, the experimental setup is described including the data sets and evaluation metrics used. In Sec. 4 the experimental results are presented and Sec. 5 gives conclusions.

## 2. Cosine-Distance VAT

In nature, the outputs of most systems are smooth w.r.t spatial and temporal inputs [9]. Prior studies have confirmed that smoothing the output distribution of a classifier (*i.e.*, encouraging the classifier to output similar distributions) against perturbations of the input can improve its generalisation performance in semi-supervised learning [1, 10, 11, 12]. In the standard version of VAT [1] (the efficacy of which has been verified by [13, 14]) an additive loss is introduced, which tries to smooth the categorical output distribution (measured by the KL-divergence) around every data point that lies on the data manifold. Here, VAT will be formulated on the level of the embedding layer rather than the output (classification) layer. The purpose of the proposed variant of VAT is to smooth the embedding generator in terms of the cosine-distance, termed cosine-distance virtual adversarial training (CD-VAT). CD-VAT should

<sup>1</sup>See General Data Protection Regulation (GDPR) or California Consumer Privacy Act (CCPA).

be used together with an angular penalty loss (such as angular softmax [15]) in comparison to the standard cross-entropy loss. These types of losses are very popular for both speaker verification [16, 17, 18] as well as face verification and identification [15, 19]. When angular penalty losses are used during speaker classification training, the resulting embedding generator produces embeddings that are angularly discriminative *i.e.* the cosine-distance between embeddings indicates if embeddings come from the same speakers.

## 2.1. CD-VAT loss

CD-VAT adds an additional loss  $\mathcal{R}_{\text{CDVAT}}$  (the CD-VAT loss) to the supervised loss  $l(\cdot, \cdot)$  with the interpolation constant  $\alpha$  and is computed on both the labelled data set  $\mathcal{D}_l$  (size  $N_l$ ) and the unlabelled data set  $\mathcal{D}_{ul}$  (size  $N_{ul}$ ). The combined loss  $\mathcal{L}$  is then used to train the parameters  $\theta$  and in turn the embedding generator  $\mathbf{e}(\mathbf{x}, \theta)$ .

$$\mathcal{L}(\mathcal{D}_l, \mathcal{D}_{ul}, \theta) = l(\mathcal{D}_l, \theta) + \alpha \mathcal{R}_{\text{CDVAT}}(\mathcal{D}_l, \mathcal{D}_{ul}, \theta) \quad (1)$$

The CD-VAT loss,  $\mathcal{R}_{\text{CDVAT}}$ , is the sum of local losses  $\text{LCS}(\mathbf{x}, \theta)$  that are computed for each input feature sequence  $\mathbf{x} \in \mathcal{D}_l, \mathcal{D}_{ul}$ .

$$\mathcal{R}_{\text{CDVAT}}(\mathcal{D}_l, \mathcal{D}_{ul}, \theta) = \frac{1}{N_l + N_{ul}} \sum_{\mathbf{x} \in \mathcal{D}_l, \mathcal{D}_{ul}} \text{LCS}(\mathbf{x}, \theta) \quad (2)$$

The local cosine smoothness,  $\text{LCS}(\mathbf{x}, \theta)$ , is calculated in two steps. First, a perturbation ( $\mathbf{r}_{\text{CDVAT}}$ ) to the input sequence  $\mathbf{x}$  is found. This perturbation is chosen to be an *adversarial*<sup>2</sup> perturbation that maximally changes the embedding ( $\mathbf{e}(\mathbf{x} + \mathbf{r}_{\text{CDVAT}}, \theta)$ ) of the input feature sequence,  $\mathbf{x}$ , as measured by the cosine-distance ( $cd[\cdot, \cdot]$ ).  $\epsilon$  is the maximum norm of  $\mathbf{r}_{\text{CDVAT}}$ .

$$\mathbf{r}_{\text{CDVAT}} = \arg \max_{\mathbf{r}; \|\mathbf{r}\| \leq \epsilon} cd\left[\mathbf{e}(\mathbf{x}, \hat{\theta}), \mathbf{e}(\mathbf{x} + \mathbf{r}, \hat{\theta})\right] \Big|_{\hat{\theta}=\theta} \quad (3)$$

$$cd[\mathbf{a}, \mathbf{b}] = \frac{1}{2} - \frac{\mathbf{a}^T \mathbf{b}}{2\|\mathbf{a}\|\|\mathbf{b}\|} \quad (4)$$

Second,  $\text{LCS}(\mathbf{x}, \theta)$  is then the cosine-distance between the embedding of the (maximally) perturbed input sequence and the embedding for the unperturbed input sequence.

$$\text{LCS}(\mathbf{x}, \theta) = cd\left[\mathbf{e}(\mathbf{x}, \hat{\theta}), \mathbf{e}(\mathbf{x} + \mathbf{r}_{\text{CDVAT}}, \hat{\theta})\right] \Big|_{\hat{\theta}=\theta} \quad (5)$$

$\hat{\theta}$  is the current setting for  $\theta$  at a particular instant during optimisation *i.e.* it is treated as a constant. The distinction between  $\theta$  and  $\hat{\theta}$  is made because gradients of  $\text{LCS}(\mathbf{x}, \theta)$  are only propagated back through the embedding generated with the input perturbation *i.e.*  $\mathbf{e}(\mathbf{x} + \mathbf{r}_{\text{CDVAT}}, \theta)$  and not  $\mathbf{e}(\mathbf{x}, \hat{\theta})$ .  $\text{LCS}(\mathbf{x}, \theta)$  indicates how ‘‘sensitive’’ the embedding of input  $\mathbf{x}$  is.

## 2.2. Approximation of $\mathbf{r}_{\text{CDVAT}}$

Given the adversarial perturbation  $\mathbf{r}_{\text{CDVAT}}$ , the optimisation of the combined loss  $\mathcal{L}$  is straightforward, because the gradients of  $\text{LCS}$  w.r.t  $\theta$  are well defined<sup>3</sup>. In this section a method for approximately finding  $\mathbf{r}_{\text{CDVAT}}$  is described. For simplicity, let:

$$cd[\mathbf{r}, \mathbf{x}, \theta] = cd\left[\mathbf{e}(\mathbf{x}, \hat{\theta}), \mathbf{e}(\mathbf{x} + \mathbf{r}, \hat{\theta})\right] \Big|_{\hat{\theta}=\theta} \quad (6)$$

$cd[\mathbf{r}, \mathbf{x}, \theta]$  has a minimum of zero at  $\mathbf{r} = \mathbf{0}$ . Therefore, the gradient w.r.t  $\mathbf{r}$  is also zero at  $\mathbf{r} = \mathbf{0}$ . Therefore, the second-order Taylor approximation of  $cd[\mathbf{r}, \mathbf{x}, \theta]$  is given by:

$$cd[\mathbf{r}, \mathbf{x}, \theta] \approx \frac{1}{2} \mathbf{r}^T \mathbf{H}(\mathbf{x}, \theta) \mathbf{r} \quad (7)$$

where  $\mathbf{H}(\mathbf{x}, \theta) = \nabla \nabla_{\mathbf{r}} cd[\mathbf{r}, \mathbf{x}, \theta] \Big|_{\mathbf{r}=\mathbf{0}}$ . For simplicity, let  $\mathbf{H} = \mathbf{H}(\mathbf{x}, \theta)$ . Under this approximation  $\mathbf{r}_{\text{CDVAT}}$  emerges as the dominant eigenvector  $\mathbf{u}(\mathbf{x}, \theta)$  of  $\mathbf{H}$  with magnitude  $\epsilon$  (see constraint from Eqn. (3)). This shows that CD-VAT in effect penalises  $\lambda_1(\mathbf{r}, \mathbf{x}, \theta)$ , the largest eigenvalue of  $\mathbf{H}$ :

$$cd[\mathbf{r}, \mathbf{x}, \theta] \approx \frac{1}{2} \epsilon^2 \lambda_1(\mathbf{r}, \mathbf{x}, \theta) \quad (8)$$

The dominant eigenvector,  $\mathbf{u}$ , of  $\mathbf{H}$  can be found using the standard power iteration method [20] combined with finite differences. Let  $\mathbf{v}_0$  be a randomly sampled vector that is not orthogonal to  $\mathbf{u}$ . Then the iterative calculation of

$$\mathbf{v}_{i+1} \leftarrow \overline{\mathbf{H}\mathbf{v}_i} \quad (9)$$

causes  $\mathbf{v}_i$  to converge to  $\mathbf{u}$ . The Hessian-vector product,  $\mathbf{H}\mathbf{v}_i$ , can be approximated based on finite differences.

$$\nabla_{\mathbf{r}} cd[\mathbf{r}, \mathbf{x}, \theta] \Big|_{\mathbf{r}=\zeta \mathbf{v}_i} \approx \nabla_{\mathbf{r}} cd[\mathbf{r}, \mathbf{x}, \theta] \Big|_{\mathbf{r}=\mathbf{0}} + \zeta \mathbf{H}\mathbf{v}_i \quad (10)$$

$$\mathbf{H}\mathbf{v}_i \approx \frac{1}{\zeta} \nabla_{\mathbf{r}} cd[\mathbf{r}, \mathbf{x}, \theta] \Big|_{\mathbf{r}=\zeta \mathbf{v}_i} \quad (11)$$

Therefore, to obtain  $\mathbf{r}_{\text{CDVAT}}$  we can use the iterative procedure:

$$\mathbf{r}_{\text{CDVAT}} \approx \epsilon \cdot \mathbf{v}_K \quad (12)$$

$$\mathbf{v}_{i+1} = \frac{\mathbf{g}_{i+1}}{\|\mathbf{g}_{i+1}\|} \quad (13)$$

$$\mathbf{g}_{i+1} = \nabla_{\mathbf{r}} cd[\mathbf{r}, \mathbf{x}, \theta] \Big|_{\mathbf{r}=\zeta \mathbf{v}_i} \quad (14)$$

where  $\mathbf{v}_K$  is the approximation of  $\mathbf{u}(\mathbf{x}, \theta)$  after  $K$  power iterations and  $\mathbf{v}_0$  is sampled uniformly on the unit-sphere. The value of  $\zeta$  should be as small as possible to get the best estimate of  $\mathbf{H}\mathbf{v}_i$ , but large enough not to cause numerical issues. Here,  $\zeta$  is set to 0.005 in all our experiments<sup>4</sup>.

$\mathbf{g}_{i+1} = \nabla_{\mathbf{r}} cd[\mathbf{r}, \mathbf{x}, \theta] \Big|_{\mathbf{r}=\zeta \mathbf{v}_i}$  is derived below:

$$\mathbf{g}_{i+1} = \frac{\partial cd[\mathbf{r}, \mathbf{x}, \theta]}{\partial \mathbf{r}} \Big|_{\mathbf{r}=\zeta \mathbf{v}_i} \quad (15)$$

$$= \frac{\partial \mathbf{e}(\mathbf{x} + \mathbf{r}, \hat{\theta})}{\partial \mathbf{r}} \cdot \frac{\partial cd[\mathbf{r}, \mathbf{x}, \theta]}{\partial \mathbf{e}(\mathbf{x} + \mathbf{r}, \hat{\theta})} \Big|_{\mathbf{r}=\zeta \mathbf{v}_i} \quad (16)$$

$$\text{let } \mathbf{e}(\mathbf{x}, \hat{\theta}) = \mathbf{e} \text{ and } \mathbf{e}(\mathbf{x} + \mathbf{r}, \hat{\theta}) = \mathbf{e}_r$$

$$\frac{\partial cd[\mathbf{r}, \mathbf{x}, \theta]}{\partial \mathbf{e}_r} = \frac{-1}{2\|\mathbf{e}\|} \cdot \frac{\partial}{\partial \mathbf{e}_r} \left( \frac{\mathbf{e}^T \mathbf{e}_r}{\|\mathbf{e}_r\|} \right) \quad (17)$$

$$= \frac{-1}{2\|\mathbf{e}\|} \cdot \left( \frac{1}{\|\mathbf{e}_r\|} \cdot \frac{\partial}{\partial \mathbf{e}_r} \left( \mathbf{e}^T \mathbf{e}_r \right) + \frac{\partial}{\partial \mathbf{e}_r} \left( \frac{1}{\|\mathbf{e}_r\|} \right) \cdot \left( \mathbf{e}^T \mathbf{e}_r \right) \right) \quad (18)$$

$$= \frac{-1}{2\|\mathbf{e}\|} \cdot \left( \frac{1}{\|\mathbf{e}_r\|} \cdot \mathbf{e} - \frac{\mathbf{e}_r}{\|\mathbf{e}_r\|^3} \cdot \left( \mathbf{e}^T \mathbf{e}_r \right) \right) \quad (19)$$

The pre-multiplication with  $\frac{\partial \mathbf{e}(\mathbf{x} + \mathbf{r}, \hat{\theta})}{\partial \mathbf{r}}$  in Eqn. (16) is equivalent to the standard back-propagation algorithm.

<sup>2</sup>Note: no relationship to generative adversarial networks (GANs).

<sup>3</sup>Eqn. (19) can be used.

<sup>4</sup>For our experiments this is a norm of 10e-6 per feature vector.

To summarise, to obtain  $\mathbf{r}_{\text{CDVAT}}$  the required calculations are:

- a forward pass to get  $\mathbf{e}(\mathbf{x}, \hat{\theta}) = \mathbf{e}$  for Eqn. (19)
- then for each power iteration:
  - a forward pass to get  $\mathbf{e}(\mathbf{x} + \mathbf{v}_i, \hat{\theta}) = \mathbf{e}_r$  for Eqn. (19)
  - a backward pass to get  $\nabla_{\mathbf{r}} cd[\mathbf{r}, \mathbf{x}, \theta] |_{\mathbf{r}=\zeta \mathbf{v}_i}$  for Eqn. (14)

Our experiments on multiple data sets suggest that  $K = 1$  is sufficient, such that  $\mathbf{r}_{\text{CDVAT}}$  does not change significantly for further iterations as measured by the dot-product of consecutive  $\mathbf{v}_i$ . This single power iteration, however, increases  $cd[\mathbf{r}, \mathbf{x}, \theta]$  by up to a factor of  $10^4$  in comparison to just using  $\mathbf{v}_0$  *i.e.* the robustness of the embedding to the simple normalised Gaussian noise  $\mathbf{v}_0$  is far larger than to the adversarial noise  $\mathbf{r}_{\text{CDVAT}}$ .

### 2.3. Related Work

Adversarial training was originally proposed by [21], where it was discovered that for image classification deep neural networks (DNNs) are very vulnerable to input perturbations applied to an input in the direction to which the model’s label assignment is the most sensitive, even when the perturbation is so small that human eyes cannot discern the perturbation. Such perturbed data points are also known as *adversarial examples* and can be used as additional data points for training [22]. For speaker verification [23] has investigated robustness to adversarial examples.

Commonly used SSL techniques in speech processing include self-training. The classifier is first trained on the labelled data and then used to assign labels to the unlabelled data. The unlabelled data (possibly after filtering) with assigned labels is combined with the labelled data. The classifier is then trained on the enlarged set of labelled data. The first successful applications of self-training were speech recognition [24, 25, 26, 27] and word-sense disambiguation [28]. For SSL of speaker embeddings self-training cannot be used if the unlabelled data does not come from the same set of speakers as the labelled data.

Self-supervised learning for speaker embeddings was used in [29] where speaker embeddings were trained via reconstructing the frames of a target speech segment, given the inferred embedding of another speech segment of the same utterance. In comparison to CD-VAT their method needs the training data to be segmented into utterances that each belong to one speaker. A purely unsupervised approach for speaker embeddings that is also based on a reconstruction loss is proposed in [30]. Though these unsupervised approaches demonstrate impressive results for generating unsupervised speaker embeddings, they still require supervised data to train the probabilistic linear discriminant analysis (PLDA) back-end used for speaker verification.

CD-VAT notably improves the intra-speaker compactness (ISC). In [31] the ISC is directly optimised by adding a supervised loss to the triplet loss that is otherwise used.

Related to VAT are consistency-based SSL methods, such as the mean teacher method [32] which was applied to audio command classification by [33]. Another popular method to reduce labelling effort is to more effectively exploit the existing labelled corpus through data augmentation [3, 34].

## 3. Experimental Setup

Experiments were designed to evaluate the effect of CD-VAT on general speaker verification performance, while also more directly testing the effect of CD-VAT on intra-speaker compactness and inter-speaker separability.

### 3.1. Data Sets

Two data sets, VoxCeleb1 (dev+test) and VoxCeleb2 (dev+test) are used to train the models and evaluate speaker verification performance. The VoxCeleb data sets consist of utterances that were obtained from Youtube videos and automatically labelled using a visual speaker recognition pipeline. The VoxCeleb2 (dev+test) data set, together with the dev portion of the VoxCeleb1 data set, is used for training. For evaluation, the test portion of VoxCeleb1 is used. The combined train set consists of more than 2750 hours of data from 7323 speakers. The evaluation set consists of 4874 utterances from 40 speakers, for which the official speaker verification list of 37720 utterance pairs is used. More information about the data is contained in Table 1.

Table 1: *The train data is the combination of VoxCeleb2 (dev+test) and the development portion of VoxCeleb1. The test data is the test portion of VoxCeleb1 for which 37720 utterance pairs are the verification list.*

Title	train	test
# Speakers	7323	40
# Videos	172299	677
# Utterances	1276888	4874
Avg. Utterance Len.	7.85 sec	8.25 sec

The system input features are 30-d mel-frequency cepstral coefficients (MFCCs). The MFCCs are extracted (using HTK [35]) using 25ms windows with 10ms frame increments from 30 filterbank channels. No vocal tract length normalization (VTLN) was applied and  $c_0$  is used instead of energy. These inputs were normalised at the utterance level for mean and globally for variance. No data augmentation was used for these experiments.

### 3.2. Model Architecture

In our model, utterance-level speaker embeddings are created by averaging multiple  $L_2$ -normalised window-level embeddings. The input window to the window-level embedding generator is around 2 seconds (213 frames, [-106,+106]) long. The shift between windows is just under 100 frames (see details below). The embedding generator uses a time-delay neural network (TDNN) [36, 37] with a total input context of [-7,+7], which is shifted from {-99} to {+99} with shifts of 6 frames (resulting in the overall input window of [-106,+106]). The 34 output vectors of the TDNN are combined using the self-attentive layer proposed in [38]. This is followed by a linear projection down to the embedding size, which is then the window-level embedding. The TDNN structure resembles the one used in the x-vector models [3] (*i.e.* TDNN-layers with the following input contexts: [-2,+2], followed by {-2,0,+2}, followed by {-3,0,+3}, followed by {0}). The first three TDNN-layers have a size of 512, the third a size of 256 and the embedding size is 32.

An utterance of length  $T$  fits  $N = \lceil \frac{T-213}{100} \rceil$  full windows (at shifts of 100 frames) plus another window if padding were used (*e.g.* replication padding to  $213 + N * 100$  frames). To avoid padding, shifts of  $\frac{T-213}{N}$  are used (*i.e.* slightly under 100 frames). The resulting indices are rounded to the nearest integer. For utterances shorter than 213 frames, the window is aligned to the centre of the utterance and replication padding used.

### 3.3. Training

CD-VAT was implemented in HTK [35] with which all models were trained in conjunction with PyHTK [39]. For training, the

window-level embedding is classified into the different speakers. The training objective for supervised training is angular softmax [15] with  $m = 1$ . The embedding generator is optimised using stochastic gradient descent (SGD) with momentum, and weightdecay was used for regularisation. The learning rate scheduler is NewBob. The batch size used for the supervised loss,  $l(\cdot, \cdot)$ , was 200 except for the model trained on the entire data set for which a batch size of 400 was used. The batch size used for the CD-VAT loss,  $\mathcal{R}_{\text{CDVAT}}$ , was 800 *i.e.* four times higher. The interpolation coefficient  $\alpha$  was set to 0.4 and the norm of the adversarial perturbation  $\epsilon$  was set to 13. The model was trained directly on the combined loss,  $\mathcal{L}(\cdot)$ , *i.e.* not pre-trained on the purely supervised loss.

For the experiments two partitions into labelled and unlabelled were created. For one 220k utterances are labelled and 440k utterances for the other. The remaining 1057k and 837k utterances, respectively, form the unlabelled dataset. The 220k and 440k utterances are chosen from the top of the utterance list sorted by official utterance name. Of the labelled dataset 20k and 40k utterances, respectively, form the validation set.

### 3.4. Evaluation Criteria (EER, ISC, ISS)

The main evaluation criterion is the speaker verification equal error rate (EER). First, utterance embeddings are formed for each utterance by averaging the window embeddings, where the windows are based on the shifts described at the end of Sec. 3.2. The scores necessary for the receiver operating characteristic (ROC) curve used in the EER calculation are the cosine-distances between the embeddings of utterance pairs, in comparison to the otherwise commonly used PLDA backend<sup>5</sup>. The ROC curve is built using scikit-learn [40].

Speaker-discriminative acoustic embeddings should have two qualities, intra-speaker compactness (ISC) *i.e.* how close to each other are the embeddings of a single speaker and inter-speaker separability (ISS) *i.e.* how far apart are the embeddings of different speakers. To give further insight about the embeddings generated from our models we attempt to give a measure of these two qualities. First all utterance embeddings belonging to the same speaker are collected and the centroid calculated (one per speaker). For the ISS the average pairwise distance between the centroids of all speakers is found. For the ISC, all utterance embeddings belonging to the same speaker are collected and the average cosine-distance to the centroid of that speaker found and the average of those per-speaker scores calculated.

## 4. Experimental Results

Tables 2 and 3 show the results of applying CD-VAT on the VoxCeleb data set. The supervised baseline model trained on 200k utterances of labelled data achieves an EER of 8.32%. The use of CD-VAT reduces the EER by 11.1% relative down to an EER of 7.40%. This represents an EER recovery of 32.5% *i.e.* we achieve 32.5% of the reduction in EER that we would get from pure supervised training if we had the speaker labels for the unlabelled part of the training data (such a model has an EER of 5.52%). This error rate recovery is similar to those seen in other areas of machine learning. For instance, [41] presents a word error rate recovery of 37% for SSL (minimum entropy training) of a DNN-HMM speech recogniser without additional language modelling data. At the same time SSL for speaker embeddings presents additional challenges as for larger numbers of speakers

<sup>5</sup>[16, 18] are other publications that experiment with direct cosine-scoring for speaker verification.

Table 2: Evaluation of CD-VAT on the VoxCeleb dataset. The evaluation criterion EER is explained in Section 3.4. For EER a lower value is better.  $\mathcal{D}_l$  is the labelled data set and  $\mathcal{D}_{ul}$  is the unlabelled data set.

System	Utts $\mathcal{D}_l$	#Speakers	Utts $\mathcal{D}_{ul}$	EER
Sup 1 CDVAT 1	200k	1249	- 1057k	8.32% 7.40%
Sup 2 CDVAT 2	400k	2504	- 837k	6.85% 6.46%
Sup 3	1277k	7323	-	5.52%

class overlap can exist. The information content of unlabelled examples decreases as classes overlap as shown by [42, 43].

The supervised baseline model trained on 400k utterances of labelled data achieves an EER of 6.85%. The use of CD-VAT reduces the EER by 5.7% relative down to an EER of 6.46%. This represents an EER recovery of 29.3%.

Table 3: Evaluation of CD-VAT on the VoxCeleb dataset. The evaluation criteria ISC and ISS are explained in Section 3.4. For ISC a lower value is better. For ISS a higher value is better.

System	Utts $\mathcal{D}_l$	Utts $\mathcal{D}_{ul}$	ISC	ISS
Sup 1 CDVAT 1	220k	- 1057k	0.13 0.09	0.38 0.36
Sup 2 CDVAT 2	440k	- 837k	0.14 0.09	0.40 0.38
Sup 3	1277k	-	0.14	0.46

Furthermore, the ISC, which is very closely related to the CD-VAT smoothing loss is reduced for the 200k and the 400k models by 31% and 36% respectively. However, at the same time the ISS is also slightly reduced. This shows one disadvantage of CD-VAT, which is that it also brings the embeddings of all utterances closer together. To put these values into perspective, the threshold of the cosine-scoring used to obtain the EER is between 0.42 and 0.48 for the systems trained.

## 5. Conclusions

We have presented cosine-distance virtual adversarial training (CD-VAT), a method that allows for semi-supervised training of speaker-discriminative acoustic embeddings without the requirement that the set of speakers is the same for the labelled and the unlabelled data. It is shown that CD-VAT can improve speaker verification performance on the VoxCeleb data set over a purely supervised baseline. The proposed method recovers 32.5% of the EER improvement that is obtained when speaker labels are available for the unlabelled data. CD-VAT also significantly improves the intra-speaker compactness (ISC) of the speaker embeddings. At the same time, however, the computational cost of CD-VAT is twice as high (per data point) as supervised training and two new hyper-parameters, that need to be tuned, are introduced.

## 6. Acknowledgements

Florian Kreyssig is funded by an EPSRC Doctoral Training Partnership Award.

## 7. References

- [1] T. Miyato, S. Maeda, S. Ishii, & M. Koyama, “Virtual adversarial training: A regularization method for supervised and semi-supervised learning,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [2] D. Snyder, D. Garcia-Romero, D. Povey, & S. Khudanpur, “Deep neural network embeddings for text-independent speaker verification,” *Proc. Interspeech*, Stockholm, 2017.
- [3] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, & S. Khudanpur, “X-vectors: Robust DNN embeddings for speaker recognition,” *Proc. ICASSP*, Calgary, 2018.
- [4] M. Diez, L. Burget, S. Wang, J. Rohdin, & J. Černocký, “Bayesian HMM based x-vector clustering for speaker diarization,” *Proc. Interspeech*, Graz, 2019.
- [5] Q. Li, F.L. Kreyszig, C. Zhang, & P.C. Woodland, “Discriminative neural clustering for speaker diarisation,” *arXiv preprint arXiv:1910.09703*, 2019.
- [6] J. Villalba, N. Chen, D. Snyder, D. Garcia-Romero, A. McCree, G. Sell, J. Borgstrom, L.P. García-Perera, F. Richardson, R. Dehak, P.A. Torres-Carrasquillo, & N. Dehak, “State-of-the-art speaker recognition with neural network embeddings in NIST SRE18 and speakers in the wild evaluations,” *Computer Speech & Language*, vol. 60, pp. 101026, 2020.
- [7] N. Dehak, P.J. Kenny, R. Dehak, P. Dumouchel, & P. Ouellet, “Front-end factor analysis for speaker verification,” *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 19, pp. 788–798, 2011.
- [8] E. Variani, X. Lei, E. McDermott, I.L. Moreno, & J. Gonzalez-Dominguez, “Deep neural networks for small footprint text-dependent speaker verification,” *Proc. ICASSP*, Florence, 2014.
- [9] G. Wahba, *Spline models for observational data*, Siam, 1990.
- [10] M. Sajjadi, M. Javanmardi, & T. Tasdizen, “Regularization with stochastic transformations and perturbations for deep semi-supervised learning,” *Proc. NIPS*, Barcelona, 2016.
- [11] S. Laine & T. Aila, “Temporal ensembling for semi-supervised learning,” *Proc. ICLR*, Toulon, 2016.
- [12] Y. Luo, J. Zhu, M. Li, Y. Ren, & B. Zhang, “Smooth neighbors on teacher graphs for semi-supervised learning,” *Proc. CVPR*, Salt Lake City, 2018.
- [13] A. Oliver, A. Odena, C. Raffel, E.D. Cubuk, & I.J. Goodfellow, “Realistic evaluation of deep semi-supervised learning algorithms,” *Proc. NeurIPS*, Montreal, 2018.
- [14] X. Zhai, A. Oliver, A. Kolesnikov, & L. Beyer, “S4: Self-supervised semi-supervised learning,” *Proc. ICCV*, Seoul, 2019.
- [15] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, & L. Song, “SphereFace: Deep hypersphere embedding for face recognition,” *Proc. CVPR*, Honolulu, 2017.
- [16] Y. Li, F. Gao, Z. Ou, & J. Sun, “Angular softmax loss for end-to-end speaker verification,” *Proc. ISCSLP*, Taipei City, 2018.
- [17] Y. Liu, L. He, & J. Liu, “Large margin softmax loss for speaker verification,” *Proc. Interspeech*, Graz, 2019.
- [18] C. Luu, P. Bell, & S. Renals, “DropClass and DropAdapt: Dropping classes for deep speaker representation learning,” *Proc. Speaker Odyssey*, Tokyo, 2020.
- [19] J. Deng, J. Guo, N. Xue, & S. Zafeiriou, “ArcFace: Additive angular margin loss for deep face recognition,” *Proc. CVPR*, Long Beach, 2019.
- [20] E. Kreyszig, H. Kreyszig, & E.J. Norminton, *Advanced engineering mathematics*, chapter 20.8 Power method for eigenvalues, pp. 885–887, Wiley, Hoboken, NJ, tenth edition, 2011.
- [21] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, & R. Fergus, “Intriguing properties of neural networks,” *Proc. ICLR*, Banff, 2014.
- [22] I.J. Goodfellow, J. Shlens, & C. Szegedy, “Explaining and harnessing adversarial examples,” *Proc. ICLR*, San Diego, 2015.
- [23] X. Li, N. Li, J. Zhong, X. Wu, X. Liu, D. Su, D. Yu, & H. Meng, “Investigating Robustness of Adversarial Samples Detection for Automatic Speaker Verification,” *Proc. Interspeech*, Shanghai (virtual), 2020.
- [24] G. Zavaliagkos & T. Colthurst, “Utilizing untranscribed training data to improve performance,” *Proc. DARPA Broadcast News Transcription and Understanding Workshop*, Landsdowne, 1998.
- [25] T. Kemp & A. Waibel, “Unsupervised training of a speech recognizer using tv broadcasts,” *Proc. ICSLP*, Sydney, Australia, 1998.
- [26] L. Lamel, J.L. Gauvain, & G. Adda, “Lightly supervised and unsupervised acoustic model training,” *Computer Speech & Language*, vol. 16, no. 1, 2002.
- [27] H.Y. Chan & P. Woodland, “Improving broadcast news transcription by lightly supervised discriminative training,” *Proc. ICASSP*, Montreal, 2004.
- [28] D. Yarowsky, “Unsupervised word sense disambiguation rivaling supervised methods,” *Proc. ACL*, Stroudsburg, 1995.
- [29] T. Stafylakis, J. Rohdin, O. Plchot, P. Mizera, & L. Burget, “Self-supervised speaker embeddings,” *Proc. Interspeech*, Graz, 2019.
- [30] Z. Peng, S. Feng, & T. Lee, “Mixture factorized auto-encoder for unsupervised hierarchical deep factorization of speech signal,” *Proc. ICASSP*, Barcelona (virtual), 2020.
- [31] N. Le & J.M. Odobez, “Robust and discriminative speaker embedding via intra-class distance variance regularization,” *Proc. Interspeech*, Hyderabad, 2018.
- [32] A. Tarvainen & H. Valpola, “Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results,” *Proc. NIPS*, 2017.
- [33] K. Lu, C.S. Foo, K.K. Teh, H.D. Tran, & V.R. Chandrasekhar, “Semi-supervised audio classification with consistency-based regularization,” *Proc. Interspeech*, Graz, 2019.
- [34] H. Yamamoto, K.A. Lee, K. Okabe, & T. Koshinaka, “Speaker augmentation and bandwidth extension for deep speaker embedding,” *Proc. Interspeech*, Graz, 2019.
- [35] S. Young, G. Evermann, M. Gales, T. Hain, D. Kershaw, X. Liu, G. Moore, J. Odell, D. Ollason, D. Povey, A. Ragni, V. Valtchev, P. Woodland, & C. Zhang, *The HTK Book*, Cambridge University Engineering Department, 2015.
- [36] V. Peddinti, D. Povey, & S. Khudanpur, “A time delay neural network architecture for efficient modeling of long temporal contexts,” *Proc. Interspeech*, Dresden, 2015.
- [37] F.L. Kreyszig, C. Zhang, & P.C. Woodland, “Improved TDNNs using deep kernels and frequency dependent grid-RNNS,” *Proc. ICASSP*, Calgary, 2018.
- [38] G. Sun, C. Zhang, & P.C. Woodland, “Speaker diarisation using 2D self-attentive combination of embeddings,” *Proc. ICASSP*, Brighton, 2019.
- [39] C. Zhang, F.L. Kreyszig, Q. Li, & P.C. Woodland, “PyHTK: Python library and ASR pipelines for HTK,” *Proc. ICASSP*, Brighton, 2019.
- [40] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, & E. Duchesnay, “Scikit-learn: Machine learning in python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [41] V. Manohar, D. Povey, & S. Khudanpur, “Semi-supervised maximum mutual information training of deep neural network acoustic models,” *Proc. Interspeech*, Dresden, Germany, 2015.
- [42] V. Castelli & T.M. Cover, “The relative value of labeled and unlabeled samples in pattern recognition with an unknown mixing parameter,” *IEEE Transactions on Information Theory*, vol. 42, no. 6, 1996.
- [43] T.J. O’Neill, “Normal discrimination with unclassified observations,” *Journal of the American Statistical Association*, vol. 73, no. 364, 1978.