

GENERATIVE ADVERSARIAL SPEAKER EMBEDDING NETWORKS FOR DOMAIN ROBUST END-TO-END SPEAKER VERIFICATION

Gautam Bhattacharya^{1,2}, Joao Monteiro², Jahangir Alam², Patrick Kenny²

¹McGill University

²Computer Research Institute of Montreal

ABSTRACT

This article presents a novel approach for learning domain-invariant speaker embeddings using Generative Adversarial Networks. The main idea is to confuse a domain discriminator so that it cannot tell if embeddings are from the source or target domains. We train several GAN variants using our proposed framework and apply them to the speaker verification task. On the challenging NIST-SRE 2016 dataset, we are able to match the performance of a strong baseline x-vector system. In contrast to the the baseline systems which are dependent on dimensionality reduction (LDA) and an external classifier (PLDA), our proposed speaker embeddings can be scored using simple cosine distance. This is achieved by optimizing our models end-to-end, using an angular margin loss function. Furthermore, we are able to significantly boost verification performance by averaging our different GAN models at the score level, achieving a relative improvement of 7.2% over the baseline.

Index Terms— GAN, Speaker Verification, End-to-End, Unsupervised Domain Adaptation

1. INTRODUCTION

Speaker Embeddings are low-dimensional vector representations that contain information relevant to a person's identity. Unlike the first generation of speaker embeddings [1], Deep Neural Networks (DNN) are able to learn robust, distributed representations directly from data. While neural network based speaker embeddings are finding applications in several downstream tasks such as speech recognition, synthesis and source separation [2, 3], they have most widely been applied to speaker verification [4, 5, 6, 7].

In the simplest case, we are given two recordings from which we can obtain the corresponding speaker embeddings. We can then obtain a verification score by computing the cosine distance (or some other simple metric) between the two embeddings. For such models to be robust, they typically need to optimize the distance metric directly. Such models are often referred to as - end-to-end, and in the context of speaker verification have proved challenging to train. Consequently,

state-of-the-art DNN speaker embeddings employ the same recipe as the popular i-vector model [8] and use LDA for dimensionality reduction and PLDA for scoring verification trials.

Verifying a speaker's identity is a challenging problem. Modern speaker verification datasets like NIST-SRE 2016 add to this challenge by introducing a mismatch between the distributions of the training and test data [9]. This phenomena is referred to as domain or covariate shift. In the case of NIST-SRE 2016, the test data consists of Cantonese and Tagalog speakers, whereas the vast majority of training speakers are talking in English.

NIST also provide a small amount of unlabelled, *in-domain*, *target* data, that can be used to compensate for the domain shift. Most domain adaptation techniques that have been proposed for speaker verification have been proposed on top of i-vectors or x-vectors. In our recent work, we proposed to learn domain-invariant speaker embeddings using domain adversarial training [10]. The proposed DANSE model is also optimized end-to-end, and produces competitive verification performance using simple cosine scoring.

In this article we extend our previous work by using Generative Adversarial Networks (GAN) to achieve unsupervised domain adaptation/invariance. We are motivated to explore this line of research given the success of GAN-based domain adaptation methods in computer vision [11, 12]. The DANSE model uses gradient reversal to achieve domain invariance, and optimizes the true minmax objective of the adversarial game [13]. Replacing gradient reversal with an explicit GAN game offers the following advantages. Firstly, GAN optimization based on an inverted label loss provides stronger gradients for learning invariant mappings than gradient reversal. Secondly, the GAN framework is more general and extendable than gradient reversal. From our analysis we show that different GAN variants produce different transformations of the feature space, and our experiments indicate that combining these feature spaces is beneficial in terms of verification performance. Both methods cast domain adaptation/invariance as an adversarial game - generate features or embeddings such that a discriminator cannot tell if they come from the source or target domain.

The nature of the adversarial game makes training GAN models inherently challenging [14]. We found that a simple way to stabilize the training of our models was to make the GANs conditional. Specifically we propose to use a modified version of the Auxiliary Classifier GAN (AuxGAN)[15]. We were able to train several GAN variants within our proposed framework, both with and without the auxiliary classifier stabilization. We find that all of the GAN models outperform the DANSE model in terms of verification performance. All the adversarial speaker embeddings proposed in this work outperform our i-vector baseline by a large margin, but are not quite able to match a state-of-the-art x-vector baseline. Interestingly, we find that we are able to beat this baseline by simply averaging the scores (cosine distances) of our different GAN models.

2. DOMAIN ADAPTATION WITH GANs

A GAN is trained through a minimax game between a generator, that maps noise vectors in the image space, and a discriminator, trained to discriminate generated images from real ones. This structure also makes GANs ideal candidates for unsupervised domain adaptation - wherein the adversarial game takes place in vector or embedding space, and the discriminators goal is to differentiate between *embeddings* from the source and target domains [11, 12].

The key component of these models is the domain discriminator (red block in Fig. 1). From our experiments we find that different discriminator configurations corresponding to several GAN variants lead to different transformations of feature space. In a standard or vanilla GAN, the discriminator is trained by optimizing the Binary Cross-Entropy Loss (BCE). The GAN game is represented by the following general form:


$$\begin{aligned} \min_D \mathcal{L}_{adv_D}(\mathbf{X}_s, \mathbf{X}_t, E) \\ \min_E \mathcal{L}_{adv_E}(\mathbf{X}_s, \mathbf{X}_t, D) \end{aligned} \quad (1)$$

Where D is the domain discriminator, E is the embedding function (generator) and $\mathbf{X}_s, \mathbf{X}_t$ represent the source and target data respectively. When training the generator, the typical approach is to use the BCE loss with inverted labels. This splits the optimization into two independent objectives, one for the generator and one for the discriminator.

Eq. 1 represents the most general form of the adversarial objective, and several adversarial domain adaptation techniques are encapsulated by this framework. By setting $\mathcal{L}_{adv_E} = -\mathcal{L}_{adv_D}$, we obtain the gradient reversal model. The advantage of using the GAN objective is that it has the same fixed-point properties as the gradient reversal but provides stronger gradients to the target mapping [11].

3. GENERATIVE ADVERSARIAL SPEAKER EMBEDDING NETWORKS

Our goal is encourage a speaker embedding model to learn domain invariant features by playing a GAN game between the feature extractor (generator) and a domain discriminator. Recently we proposed the DANSE model for speaker verification, which integrates adversarial training with a speaker embedding model. We were able to show that by combining verification with domain adversarial training, we are able to achieve competitive performance using simple cosine scoring. In this work we choose to replace the gradient reversal method with an explicit GAN game. We show that the framework is robust, with several GAN variants displaying good speaker verification performance.



While the GAN game encourages domain invariance, we also require the embeddings to be speaker discriminative. These characteristics are induced by minimizing an appropriate task loss:

$$\begin{aligned} \min_{E, C} \mathcal{L}_{task}(\mathbf{X}_s, \mathbf{Y}_s) = \\ - \mathbb{E}_{(x_s, y_s) \sim (X_s, Y_s)} \sum_{k=1}^K \mathbb{1}_{[k=y_s]} \log(C(E^s(x_s))) \end{aligned} \quad (2)$$

Where $C(\cdot)$ is a classifier and $E(\cdot)$ is the Embedding function. We use the Additive Margin Softmax loss for \mathcal{L}_{task} , which directly optimizes for cosine similarity between classes [16]. It offers a clear advantage over standard cross-entropy training [10]. The AM-Softmax loss maintains the same structural form as the cross-entropy loss:

$$\begin{aligned} \mathcal{L}_{task} &= -\frac{1}{n} \sum_{i=1}^n \log \frac{e^{s \cdot (\cos \theta_{y_i} - m)}}{e^{s \cdot \cos(\theta_{y_i} - m)} + \sum_{j \neq y_i} e^{s \cdot (\cos \theta_j)}} \\ &= -\frac{1}{n} \sum_{i=1}^n \log \frac{e^{s \cdot (W^T \mathbf{f}_i - m)}}{e^{s \cdot (W^T \mathbf{f}_i - m)} + \sum_{j \neq y_i} e^{s \cdot (W^T \mathbf{f}_j)}} \end{aligned} \quad (3)$$

Where s is a scale factor and m is a margin. During training, the first step is to update the $C(\cdot)$ and $E(\cdot)$ (blue and green blocks in Fig. 1) by minimizing \mathcal{L}_{task} . Next we extract

source and target embeddings for GAN training. We first update the domain discriminator $D(\cdot)$ using BCE loss on both source and target data:

$$\begin{aligned}\mathcal{L}_{adv_D}(\mathbf{X}_s, \mathbf{X}_t, E) = & \\ & - \mathbb{E}_{x_s \sim X_s} [\log(D(E(x_s)))] \\ & - \mathbb{E}_{x_t \sim X_t} [\log(1 - D(E(x_t)))]\end{aligned}\quad (4)$$

Finally, we update the generator/embedding function, $E(\cdot)$ to fool the discriminator using the inverted label loss:

$$\mathcal{L}_{adv_E}(\mathbf{X}_s, \mathbf{X}_t, D) = -\mathbb{E}_{x_s \sim X_s} [\log(D(E(x_t)))] \quad (5)$$

Note that the embedding function $E(\cdot)$ gets updated twice, first with the task loss followed by the adversarial loss.

3.1. Auxiliary Classifier GAN

The Auxiliary Classifier GAN (AuxGAN) model augments the standard GAN framework with an auxiliary loss to perform conditional image generation [15]. This approach aims to predict side information (such as class labels), as opposed to feeding the same information to the generator and discriminator. In the context of this work, we use the auxiliary loss for regularization and representation learning.

$$\begin{aligned}\min_D \mathcal{L}_{adv_D}(\mathbf{X}_s, \mathbf{X}_t, E) + \mathcal{L}_{Aux}(\mathbf{X}_s, Y_s) \\ \min_E \mathcal{L}_{adv_E}(\mathbf{X}_s, \mathbf{X}_t, D) + \mathcal{L}_{Aux}(\mathbf{X}_s, Y_s)\end{aligned}\quad (6)$$

Eq. (6) is a modified version of the AuxGAN objective. In particular, the original formulation also uses the auxiliary loss to train on fake data as well (with fake data being assigned its own unique label). We use the cross-entropy loss over speakers as the auxiliary loss. In our experiments we found that using an auxiliary classifier helps stabilize all the GAN variants, but does not always lead to the better verification performance.

3.2. GAN Variants

Since their introduction, GANs have been one of the most researched topics in the deep learning community. Several variations of the original formulation have been proposed, each with different generative characteristics and stability issues. In this work we explore two GAN variants in addition to the standard GAN - Least-Squares GAN and Relativistic GAN [17, 18]. The former takes advantage of a well known shortcoming of the least-squares loss and forces the encoder to learn embeddings that are closer to the decision boundary. The latter model encourages true (source) data to be classified as fake (target), which we hypothesize helps to learn a more invariant feature space.

4. EXPERIMENTS & RESULTS

Training Data (Source): We used audio from previous NIST-SRE evaluations (2004-2010) and Switchboard Cellular audio for training the proposed DANSE model as well as the x-vector and i-vector baseline systems. We also augment our data with noise and reverberation, as in [4]. We add 128k noisy copies to the clean speech, ending up with 220k recordings in our training set. For training adversarial models we filter out speakers with less than 5 recordings, ending up with approximately 6000 speakers. Whereas the x-vector and i-vector systems were trained using the Kaldi recipe. We note that the performance of the baseline systems can be improved by adapting the PLDA classifier to the target data [19].

Model: In order to make a fair comparison, we use an identical network to the DANSE model. The *Embedding function/Generator*, $E(\cdot)$, consists of a 3×23 input convolutional layer, 4 residual blocks [3,4,6,3], an attentive statistics layer and two fully connected layers (512,512). The *classifier*, $C(\cdot)$, module consists of a fully connected layer (64) and the AM-Softmax output layer. The former is the final domain invariant speaker embedding extracted during evaluation. Finally, the *domain discriminator* module consists of two fully connected layers (256,256) and a binary cross-entropy output layer. Exponential Linear Units (ELU) are used as non-linear activations for all layers of the network. Batch Normalization is used on all layers except the attentive statistics layer. The s and m hyper-parameters of the AM-Softmax loss are set to 30 and 0.6 respectively. We refer the reader to [10] for a detailed description of the model.

Optimization: We start by pre-training the Embedding function using standard cross-entropy training. Pre-training is carried out using the RMSprop optimizer with a learning rate (lr) of 0.001. For training GAN based speaker embedding models we use different optimizers for training the three networks (Embedding function, Classifier, Discriminator). The classifier is optimized using RMSprop with $lr=0.003$, while the domain classifier and feature extractor are trained using SGD with $lr=0.001$. We were able to train all our GAN models using the same set of hyper-parameters. We used performance on held out validation set to determine when to stop training.

Data Sampling: We use an extremely simple approach for sampling data during training. We sample random chunks of audio (3-8 seconds) from each recording in the training set. We sample each recording 10 times to define an epoch. For each mini-batch of source data, we randomly sample (with repetition) a mini-batch from the unlabelled adaptation data for GAN training.

Speaker Verification: At test time we discard the domain discriminator branch of the model, as it is not needed for extracting embeddings. Extraction is done by performing a forward pass on the full recording, and using the 64-dimensional

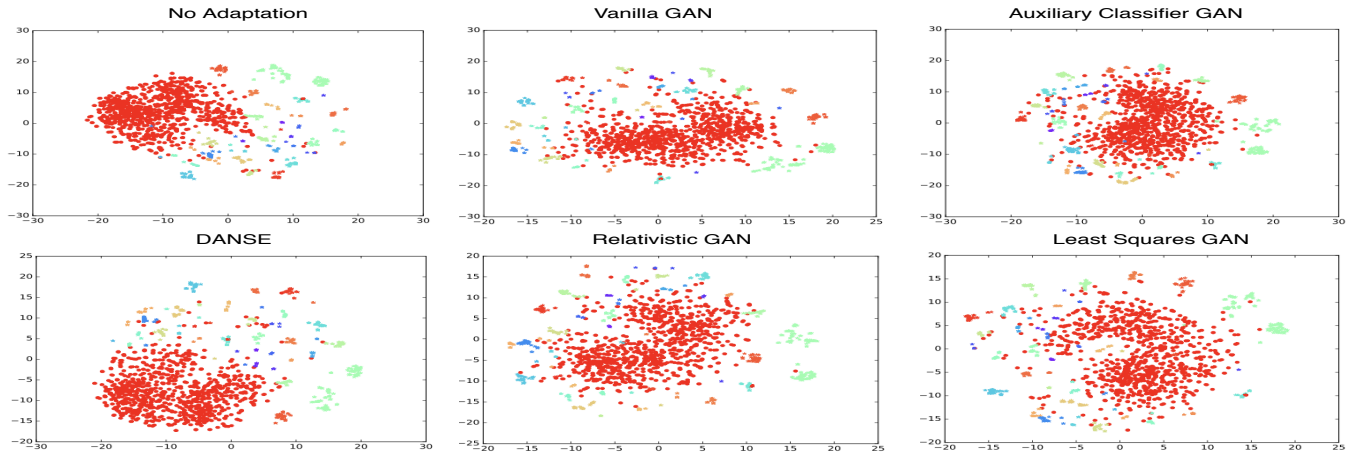


Fig. 1: t-SNE Visualization [20] of Embedding Space. Large red cluster represents target data. Other colours are source domain speakers.

final hidden layer as our speaker embedding. Verification trials are scored using cosine distance. Verification performance is reported in terms of Equal Error Rate (EER).

NIST-SRE 2016: Unlike previous years, The 2016 edition of the NIST-SRE introduced a challenging new dataset containing Cantonese and Tagalog speakers. We use the Kaldi recipes for our baseline i-vector and x-vector systems. We note this x-vector baseline may be considered as state-of-the-art performance on this dataset.

Adaptation Data (Target): 2272 unlabelled, target data recordings are provided to adapt verification systems.

<i>Model</i>	<i>Classifier</i>	<i>Cantonese</i>	<i>Tagalog</i>	<i>Pooled</i>
i-vector	LDA/PLDA	9.51	17.61	13.65
x-vector	COSINE	36.44	41.07	38.69
x-vector	LDA/PLDA	7.52	15.96	11.73
x-vector	PLDA	7.99	18.46	13.32

Table 1: Baseline Systems

<i>Model</i>	<i>Classifier</i>	<i>Cantonese</i>	<i>Tagalog</i>	<i>Pooled</i>
DANSE	COSINE	8.84	17.87	13.29
SGAN	COSINE	8.32	17.51	12.93
AuxGAN	COSINE	7.96	15.90	11.93
LSGAN	COSINE	7.90	15.63	11.74
RelGAN	COSINE	8.01	16.22	12.21
FuseGAN	COSINE	6.93	14.84	10.88

Table 2: **DANSE:** Gradient Reversal, **SGAN:** standard, **AuxGAN:** auxiliary classifier, **LSGAN:** least squares, **RelGAN:** relativistic, **FuseGAN:** score averaging

Tables 1. & 2. compare the performance of the proposed adversarial speaker embeddings against the baseline systems. Among the baseline systems we see that the DNN based x-

vector system clearly outperforms i-vectors, however only with the addition of LDA dimensionality reduction.

We see that all the GAN based models outperform DANSE by a large margin. This result confirms the hypothesis that the GAN game provides a more effective learning signal than the gradient reversal framework. All the individual GAN models also outperform the i-vector baseline by a large margin, and the best model (LSGAN) is able to match the x-vector baseline. We find that we are able to significantly improve on this baseline by simply averaging the verification scores (cosine distances) of the AuxGAN, LSGAN and RelGAN embeddings. The FuseGAN ensemble achieves the lowest EER on all three splits of the verification trials.

This result is interesting as it suggests that the different GAN models cover different modes of the target data distribution, and capture speaker information that is different yet complementary. Comparing the transformed feature spaces in Fig. 2, we see that the DANSE model appears to mainly rotate the feature space, whereas the transformations induced by the GAN models is more pronounced. Importantly we see that the source data speaker clusters remain intact after adversarial adaptation.

5. CONCLUSION

We presented a novel framework for learning domain invariant speaker embeddings using GANs. By learning several different GAN variants and combining them at the score level, we are able to achieve state-of-the-art verification performance. We find that models that explicitly encourage domain confusion through a least-squares or relativistic loss are the best suited for this task. In future work we will consider other adversarial strategies like combining feature space and data space GANs [21, 22].

6. REFERENCES

- [1] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.
- [2] Y. Jia, Y. Zhang, R. J. Weiss, Q. Wang, J. Shen, F. Ren, Z. Chen, P. Nguyen, R. Pang, I. L. Moreno *et al.*, "Transfer learning from speaker verification to multispeaker text-to-speech synthesis," *arXiv preprint arXiv:1806.04558*, 2018.
- [3] Q. Wang, H. Muckenhirn, K. Wilson, P. Sridhar, Z. Wu, J. Hershey, R. A. Saurous, R. J. Weiss, Y. Jia, and I. L. Moreno, "Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking," *arXiv preprint arXiv:1810.04826*, 2018.
- [4] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," *Submitted to ICASSP*, 2018.
- [5] C. Li, X. Ma, B. Jiang, X. Li, X. Zhang, X. Liu, Y. Cao, A. Kannan, and Z. Zhu, "Deep speaker: an end-to-end neural speaker embedding system," *arXiv preprint arXiv:1705.02304*, 2017.
- [6] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4879–4883.
- [7] G. Bhattacharya, M. J. Alam, V. Gupta, and P. Kenny, "Deeply fused speaker embeddings for text-independent speaker verification," *Proc. Interspeech 2018*, pp. 3588–3592, 2018.
- [8] P. Kenny, "Bayesian speaker verification with heavy-tailed priors," *Proc. Speaker Odyssey 2010*, pp. 3588–3592, 2010.
- [9] S. O. Sadjadi, T. Kheyrkhah, A. Tong, C. Greenberg, E. S. Reynolds, L. Mason, and J. Hernandez-Cordero, "The 2016 nist speaker recognition evaluation," in *Proc. Interspeech*, 2017, pp. 1353–1357.
- [10] G. Bhattacharya, J. Alam, and P. Kenny, "Adapting end-to-end neural speaker verification to new languages and recording conditions with adversarial training," in *Acoustics Speech and Signal Processing (ICASSP), 2019 IEEE International Conference on*. IEEE, 2019.
- [11] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *Computer Vision and Pattern Recognition (CVPR)*, vol. 1, no. 2, 2017, p. 4.
- [12] J. Shen, Y. Qu, W. Zhang, and Y. Yu, "Wasserstein distance guided representation learning for domain adaptation," in *AAAI*, 2018.
- [13] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *arXiv preprint arXiv:1409.7495*, 2014.
- [14] L. Mescheder, A. Geiger, and S. Nowozin, "Which training methods for gans do actually converge?," *International Conference on Machine Learning*, pp. 3478–3487, 2018.
- [15] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier gans," *arXiv preprint arXiv:1610.09585*, 2016.
- [16] W. Feng, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face verification," *IEEE Signal Processing Letters* 25, no. 7, pp. 926–930, 2018.
- [17] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *Computer Vision (ICCV), 2017 IEEE International Conference on*. IEEE, 2017, pp. 2813–2821.
- [18] A. Jolicoeur-Martineau, "The relativistic discriminator: a key element missing from standard gan," *arXiv preprint arXiv:1807.00734*, 2018.
- [19] D. Garcia-Romero, A. McCree, S. Shum, N. Brummer, and C. Vaquero, "Unsupervised domain adaptation for i-vector speaker recognition," in *Proceedings of Odyssey: The Speaker and Language Recognition Workshop*, 2014.
- [20] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research* 9, no. (2008), pp. 2579–2605, 2008.
- [21] S. Sankaranarayanan, Y. Balaji, C. D. Castillo, and R. Chellappa, "Generate to adapt: Aligning domains using generative adversarial networks," *ArXiv e-prints, abs/1704.01705*, 2017.
- [22] J. Hoffman, E. Tzeng, J.-Y. Zhu, P. Isola, K. Saenko, A. A. Efros, and T. Darrell, "Cycada: Cycle-consistent adversarial domain adaptation," *ArXiv e-prints, abs/1711.03213*, 2017.