

# Prosodic-Enhanced Siamese Convolutional Neural Networks for Cross-Device Text-Independent Speaker Verification

Sobhan Soleymani, Ali Dabouei, Seyed Mehdi Iranmanesh, Hadi Kazemi,  
Jeremy Dawson, and Nasser M. Nasrabadi, *Fellow, IEEE*  
West Virginia University

{ssoleyma, ad0046, seiranmanesh, hakazemi}@mix.wvu.edu,  
{jeremy.dawson, nasser.nasrabadi}@mail.wvu.edu

## Abstract

*In this paper a novel cross-device text-independent speaker verification architecture is proposed. Majority of the state-of-the-art deep architectures that are used for speaker verification tasks consider Mel-frequency cepstral coefficients. In contrast, our proposed Siamese convolutional neural network architecture uses Mel-frequency spectrogram coefficients to benefit from the dependency of the adjacent spectro-temporal features. Moreover, although spectro-temporal features have proved to be highly reliable in speaker verification models, they only represent some aspects of short-term acoustic level traits of the speaker's voice. However, the human voice consists of several linguistic levels such as acoustic, lexicon, prosody, and phonetics, that can be utilized in speaker verification models. To compensate for these inherited shortcomings in spectro-temporal features, we propose to enhance the proposed Siamese convolutional neural network architecture by deploying a multilayer perceptron network to incorporate the prosodic, jitter, and shimmer features. The proposed end-to-end verification architecture performs feature extraction and verification simultaneously. This proposed architecture displays significant improvement over classical signal processing approaches and deep algorithms for forensic cross-device speaker verification.*

## 1. Introduction

Speech is considered as a form of biometric verification, since everybody has his or her unique voice. Speaker verification aims to extract features from a speaker's speech samples and use them to recognize or verify speaker identity through modelings of the speaker's speech samples [22]. The speaker-verification literature focuses on designing a setup in which the claimed identity of a speaker is either accepted or rejected, which can be conducted as text-

dependent [7, 32, 20] or text-independent [15, 33]. During text-dependent speaker verification the speech content is a predefined, fixed text, such as a passphrase, while text-independent speaker verification aims to verify the speaker using freeform spoken words, independent of the text or language or other prior constraints. The possible unconstrained variations in text-independent speaker verification make it much more challenging compared to text-dependent models [32].

Voice samples can be acquired through different recording devices and are subject to device and quality mismatch. In addition, the samples can be recorded at different sampling rates and distances, which result in bit-rate mismatch and channel noise. The samples are also subject to background noise problem due to environmental noise and distortion. Channel-independent speaker verification frameworks [25] try to address this problem. Channel-independent text-independent frameworks are considered to be the ultimate test in the speaker verification domain [26, 14].

Deep learning algorithms are the state-of-the-art frameworks for many biometric applications such as face [17], fingerprint [9], and iris [30] classification, as well as multimodal classification [30, 31], attribute-enhanced classification [18], and domain adaptation [24]. Deep learning architectures have recently proven to be able to provide superior performance compared to traditional speaker verification algorithms, showing significant gains over the state-of-the-art Gaussian Mixture Models and Hidden Markov Models [28, 21, 23]. The majority of the deep learning architectures proposed for speaker recognition task are multilayer perceptron (MLP)-based models using Mel-frequency cepstral coefficients (MFCCs) [29, 7, 32]. However, MLP-MFCC architectures fail to preserve the correlation between the adjacent features. To address this issue convolutional neural networks (CNNs) are used in speaker recognition [33, 25]. Additionally, compared to architectures

requiring hand-crafted features, convolutional neural networks (CNNs) extract and classify features simultaneously, and, therefore, avoid losing valuable information [19].

State-of-the-art deep speaker recognition systems use spectro-temporal voice features [25, 29]. The most well-known these short-term features used in the literature are spectrogram, MFCCs, and Mel-frequency spectrogram coefficients (MFSCs). Inheriting from the short-term nature of these features, most of the models proposed only explore the acoustic level of the signal, such as spectral magnitudes and formant frequencies [5, 13]. However, several important linguistic levels such as lexicon, prosody or phonetics cannot be recognized from short-term features. These levels of information are learned habits by the speaker. These features do not perform as well as the short-term features in the identification and verification scenarios when the utterances are significantly short. However, when the length of the utterance increases, it is shown that the identification and verification performance of the prosodic features increases drastically [27]. These features also significantly improve the model when fused with short-term features [13, 12].

In the speaker-verification literature a three-phase procedure is defined. Initially, in the *development* phase, background models are developed from a large collection of data. New speakers are added to the model during the *enrollment* phase to construct speaker-dependent models. In the *evaluation* phase test utterances are compared to the enrolled speaker models and the background model to verify the identity of the speaker [32]. In this setup, the difference between low-dimensional representations of enrollment and test utterances is considered to accept or reject the hypothesis [29]. However, in the proposed algorithm, the enrollment phase is excluded. The proposed Siamese model is trained using the utterances from the training set in the *training* phase. In the *test* phase, the trained model is deployed to compute the distance between two utterances. The computed inter sub-network distance is used to determine whether or not the utterances belong to the same speaker.

In this paper, we make the following contributions: (i) prosodic, jitter, and shimmer features are deployed to enhance the performance of the proposed CNN Siamese network; (ii) a text-independent embedding space is constructed considering short-term and prosodic features; (iii) rather than extracting the features using hand-crafted methods, a fully data-driven architecture using fused CNN and MLP networks has been optimized for joint domain-specific feature extraction and representation with the application of speaker verification, finally (iv) the proposed algorithm can be used for real-time applications since it does not require the enrollment phase.

## 2. Prosodic features to enhance deep coupled CNN

CNN architectures have recently proven to outperform the traditional speaker verification algorithms. Following the scenario deployed in image processing literature, the input fed into the CNN is a nonlinearly scaled spectrogram with its first and second temporal derivatives [3]. CNN models prefer inputs that change smoothly along both dimensions. Therefore, acoustic features need to smoothly change both in time and frequency [4]. Since the acoustic signal is smooth in time, the frequency features need to preserve the locality of the speech signal. The majority of the works using deep neural networks for speech processing use MFCCs [29, 32, 6].

However, these features do not preserve the locality of the frequency domain signal since the discrete cosine transform (DCT) projects the spectral energies into a basis that does not maintain locality [3]. Recently, MFSCs have been introduced to compensate for this shortcoming [4]. MFSCs are the log-energy computed directly from the mel-frequency spectral coefficients, which are the representation of the smoothed spectral envelope of the speech. These features, which are computed similarly to MFCC features with no DCT operation, along with their deltas and delta-deltas (first and second temporal derivatives) are fed into CNN as three channels of the input, describing the acoustic energy distribution of the spoken utterances.

These short-term coefficients represent the spectral envelope of a speech frame. Although these parameters are speaker specific, they are unable to represent supra-segmental characteristics of the speech signal [13]. On the other hand, prosodic coefficients represent features that are larger than phonetic units such as; sound, duration, tone and intensity variation.

Although within-speaker variability in phonetic content and speaking style degrades the performance of speaker verification systems for short utterances [27], due to the practical complexity of the CNN architecture and the vast number of parameters that need to be trained, it is not feasible to feed the utterances to the network since it will drastically reduce the number of samples in the training set. To compensate for this shortcoming, we propose to compute the prosodic features from the whole utterances. For each utterance, several short utterances are randomly chosen. Each of these short utterances, along with the prosodic features calculated for the utterance, are fed to the network. The decision is made upon the computed overall scores.

Following the setup in [13], 18 prosodic features are extracted from the utterances: three features related to word and segmental durations (number of frames per word and length of word-internal voiced and unvoiced segments), six features related to fundamental frequency (mean, maxi-

mum, minimum, range, pseudo-slope and slope), and nine jitter and shimmer measurements. Jitter indices used in this setup are absolute jitter, relative jitter, rap, and ppq5, while the shimmer indices used are shimmer (dB), relative shimmer, apq3, apq5, and apq11 [13].

Jitter and shimmer are defined as the indices for the cycle-to-cycle variations of fundamental frequency and amplitude, respectively. These indices are used to describe the voice quality. The frequency of a speaker’s voice varies from one cycle to the next cycle. Jitter is defined as the cycle-to-cycle variation of fundamental frequency, and is the measurement of vocal stability. On the other hand, Shimmer is the index for vocal amplitude perturbation. Since these features characterize particular voices, they provide speaker-specific information.

### 3. Proposed speaker-verification architecture

The proposed Siamese architecture consists of two sub-networks that share weights. Each sub-network includes MLP and CNN networks, and the joint representation layer. Segmental features are extracted from each utterance, and are fed to the MLP network, while random short utterances are chosen from the utterances. MFSCs are extracted from short utterances and fed into CNN network. Each sub-network is represented by a fully-connected fusion layer that act as the joint representation. Two joint representations are used to train the network through the contrastive loss.

#### 3.1. Frequency- and prosody-domain networks

Pooling algorithms are used in CNN architectures to reduce the possibility of over-fitting. *Maxpooling* is the sample-based conventional process in CNN architectures to down-sample the feature map representation without smoothing the feature maps, while extracting the most important features. It reduces the maps’ dimensionality and allows the architecture to make sub-region assumptions. However, maxpooling is a shift-invariant operator and risks undesirable phonetic confusion [11]. To compensate, we propose to use multiple maxpooling sizes in the frequency domain instead of the conventional maxpooling and concatenate the output feature maps in depth as shown in Table 1. On the other hand, since the proposed architecture is text-independent, the conventional last-layer fully-connected layer is replaced by average pooling along the time axis and a fully connected layer along the frequency axis. Additionally, this modification allows the inputs to vary in size in the time domain.

The frequency domain network is comprised of five major convolutional components and two fully-connected layers which are connected in series. Each convolutional layer is followed by a rectified linear unit (ReLU) layer and a time domain maxpooling. *Conv2* and *Conv4* are also fol-

Table 1: The MFSC-dedicated CNN architectures. Notations (t) and (f) represent time and MFSC axis respectively.

CNN-MFSC			
layer	kernel	input	output
conv1	$3 \times 3 \times 64$	$40 \times 300 \times 3$	$40 \times 300 \times 64$
maxpool1 (t)	$2 \times 1 \times 1$	$40 \times 300 \times 64$	$40 \times 150 \times 64$
conv2	$3 \times 3 \times 128$	$40 \times 150 \times 64$	$40 \times 150 \times 128$
maxpool2 (t)	$2 \times 1$	$40 \times 150 \times 128$	$40 \times 75 \times 128$
maxpool2-a (f)	$1 \times 2 \times 1$	$40 \times 75 \times 128$	$20 \times 75 \times 128$
maxpool2-b (f)	$1 \times 3 \times 1$	$40 \times 75 \times 128$	$20 \times 75 \times 128$
maxpool2-c (f)	$1 \times 4 \times 1$	$40 \times 75 \times 128$	$20 \times 75 \times 128$
conv3	$3 \times 3 \times 256$	$20 \times 75 \times 384$	$20 \times 75 \times 256$
maxpool3 (t)	$2 \times 1 \times 1$	$20 \times 75 \times 256$	$20 \times 37 \times 256$
conv4	$3 \times 3 \times 256$	$20 \times 37 \times 256$	$20 \times 37 \times 256$
maxpool4	$2 \times 1 \times 1$	$20 \times 27 \times 256$	$20 \times 18 \times 256$
maxpool4-a (f)	$1 \times 2 \times 1$	$20 \times 18 \times 256$	$10 \times 18 \times 256$
maxpool4-b (f)	$1 \times 3 \times 1$	$20 \times 18 \times 256$	$10 \times 18 \times 256$
maxpool4-c (f)	$1 \times 4 \times 1$	$20 \times 18 \times 256$	$10 \times 18 \times 256$
conv5	$3 \times 3 \times 512$	$10 \times 18 \times 768$	$10 \times 18 \times 512$
avgpool (t)	$1 \times 18 \times 1$	$10 \times 18 \times 512$	$10 \times 1 \times 512$
FC6 (f)	$10 \times 1 \times 1024$	$10 \times 1 \times 512$	$1 \times 1 \times 1024$
FC7	$1 \times 1 \times 128$	$1 \times 1 \times 1024$	$1 \times 1 \times 128$

lowed by a heterogeneous frequency domain maxpooling. In the proposed heterogeneous maxpooling, different kernel sizes are applied on the feature maps and the outputs are concatenated in depth and fed into the next convolutional layer. The inputs to the frequency-domain network represent short-term features of the acoustic signal. 18 Prosodic features are fed into a multilayer perceptron with two hidden layers. Each hidden layer consists of 64 hidden units, while the output layer includes 32 nodes.

#### 3.2. Speaker-verification coupled CNN

The final objective of the proposed model is to verify whether or not two utterances recorded on different devices belong to the same speaker or not. The utterances can also be recorded at the same time or in different sessions. Therefore, the proposed method needs to satisfy the text-independent condition. On the other hand, it is not feasible to feed the whole utterances to the network, since it drastically reduces the number of samples. In addition, since the utterance can vary in length, feeding them to the network limits the batch normalization benefits. Therefore, we propose to randomly choose several fixed-length short utterances. Each short utterance is fed into the network along with the prosodic features calculated from the long utterance. The final decision is made upon the distances (scores) given to each pair of short utterances.



Training the network as classifier facilitates the extraction of the discriminative features from MFSC coefficients.

The inputs are 3 seconds utterances which are represented as  $300 \times 224 \times 3$  images. Three channels represent static, delta and delta-delta feature maps, while 40 represent the number of MFSC coefficients. The training algorithm is deployed by minimizing the softmax cross-entropy loss using mini-batch stochastic gradient descent with momentum. The training was regularized by weight decay and 50% dropout for the fully connected layers except for the last layer. The batch size, momentum and  $L_2$  penalty multiplier are set to 32, 0.9 and 0.0005, respectively. The initial learning rate is set to 0.1. The learning rate is decreased exponentially by a factor of 0.1 for every 2 epochs of training. In this network, batch normalization [16] is applied. The moving average decay is set to 0.99.

Similarly, the MLP network is optimized independently. The parameters for this optimization are the same as the parameters for the CNN network. To train the joint representation, the CNN and MLP networks are frozen and the joint representation layer is optimized greedily upon the extracted features. The initial learning rate is reduced to the smallest final learning rate among two networks. Finally, the classification architecture is trained jointly.

To train the Siamese network, the network is initialized with the weights optimized for the classifier network. The pairs are fed into the network and the contrastive loss function is minimized while the sub-networks share weight.

## 4.2. Hyperparameter optimization

The hyperparameters in our experiments are :  $\lambda$  the regularization parameter,  $\alpha_0$  initial learning rate,  $n$  number of epochs per decay for the learning rate,  $d$  moving average decay, and  $m$  as the momentum. For each optimization, the 5-fold cross-validation method on the training set is used to estimate the best hyperparameters.

## 5. Experiments and discussions

### 5.1. Dataset

**FBI Voice Collection 2016:** This database consists of two sessions (July 2016 and January 2017) of speech from 411 individuals using three recording devices: a high-quality microphone, a typical interview room recording system/DVR, and a digital recorder capturing the speech over a cell phone connection. The last two recording are recorded simultaneously. The number of male and female speakers are 205 and 206 respectively. This database is one of the few databases that allows disjoint channel-independent training and testing of the proposed algorithm. The total number of utterances is equal to 2,418. The training is conducted on 361 speakers. The test is performed on the remaining 50 subjects. A summary of the database is presented in Table 2.

Table 2: The number of utterances in train and test sets for FBI voice collection 2016 database.

	Train set	Test set
Utterances	2148	300
Microphone	705	100
DVR	708	100
Phone	705	100

### 5.2. Data representation

Initially all utterances are re-sampled to 48 KHz. For each utterance prosodic features were extracted using Praat software for acoustic analysis [1]. These 18 features listed in Section 2 are the inputs fed to MLP network. Then, voiced segments of the utterances are detected using the voicebox toolbox [2]. Each voiced utterance is divided into 25ms frames with 60% overlap.

Each frame is multiplied with a hamming window to keep the continuity of the first and the last points in the frame. 40 MFSC coefficients are extracted from each frame. Delta and delta-delta channels are constructed for each frame as the first derivative and second temporal derivative of MFSC features. Cepstral mean and variance normalization are applied on each utterance, in which each frequency bin is normalized to zero mean and unit variance. Finally, short utterances of three seconds length with two seconds overlap are generated. The inputs fed to CNN network are  $40 \times 300 \times 3$  short utterances.

### 5.3. Training and test phases

**Training phase:** Pairs of short utterances are randomly chosen, while we make sure that the overall number of genuine and imposter pairs are equal. The pairs of short utterances are fed into the architecture along with the prosodic features. The architecture is trained under contrastive loss with no normalization on the last fully-connected layer (FC8). Here the short utterances are assumed to be independent samples, and the contrastive loss is applied on each pair of short utterances. The contrastive loss margin is set to 10.

**Test phase:** For each pair of utterances, 500 pairs of short utterances are randomly chosen. The pairs of short utterances are fed into the architecture along with the prosodic features. For each pair of short utterances, the distance is computed as the Euclidean distance between the samples in the embedding space. The vector of the distances between the short utterances is used to determine the distance between two utterances. The short samples can be noisy or may not include speaker specific information. Therefore,

Table 3: Verification performance on FBI Voice Collection 2016 database.

Algorithm	AUC	EER
i-vector/PLDA (MFSC)	0.9153	0.1579
i-vector/PLDA (MFCC)	0.9185	0.1526
Chen et al. [7]	0.9207	0.1451
Nagrani et al. [25]	0.9215	0.1469
CNN network	0.9218	0.1421
Prosodic network	0.9011	0.1673
Score-level fusion	0.9148	0.1578
Coupled network	0.9358	0.1311

averaging the distances between pairs of short utterances may include outliers. To remove the effect of these short utterances, the vector’s mean and standard deviation are computed. The average of the elements in the vicinity of two standard deviations from the mean value represent the distance between the pair of utterances.

#### 5.4. Evaluation metrics

The performance of different experiments are reported and compared using two verification metrics. The utilized metrics are equal error rate (EER) and area under curve (AUC). When false acceptance and false rejection rate for the model are equal, the common value is referred to as EER. AUC represents the area under the receiver operating characteristic curve.

#### 5.5. Results

Table 3 presents the verification results for the proposed algorithm. In addition, the verification results for CNN and MLP trained independently are presented. The score-level fusion of two networks is considered as well. The performance of the proposed algorithm is compared with that of i-vector/PLDA algorithm [10]. The same MFSC feature used in the proposed deep algorithm are used in i-vector algorithm. The i-vector model is also trained with MFCC features. The algorithm is also compared with two state-of-the-art deep architectures [25, 7].

Table 4 presents the results for channel-dependent setup. In this special case, each sub-network is fed with utterances from a specific device. To train this architecture, the sub-networks do not share weights. To initialize the parameters for this setup, both the sub-networks are initialized with the parameters from channel-independent setup. This setup leads to better performance compared to channel-independent setup, since the channel-dependent information in the test samples can be learned during the training phase. The only exception is Phone-DVR cross-device verification setup, where, both devices are considered low-quality devices.

Table 4: EER and AUC for the channel-dependent setup.

Device	Microphone	DVR	Phone
Microphone	0.0712	0.1132	0.1247
DVR		0.0827	0.2069
Phone			0.1316

(a) EER

Device	Microphone	DVR	Phone
Microphone	0.9785	0.9537	0.9512
DVR		0.9717	0.8467
Phone			0.9547

(b) AUC

## 6. Conclusion

In this paper we proposed a novel cross-device text-independent speaker verification Siamese architecture, where Mel-frequency spectrogram coefficients are used to benefit from correlation of the adjacent features. In addition, prosodic features were deployed to enhance the spectral features fed to CNN. A MLP network is trained to represent the prosodic features describing words, fundamental frequency, jitter and shimmer. The joint representation fusing two networks, trains the network through contrastive loss. The proposed end-to-end verification architecture performs feature extraction and verification simultaneously. The proposed architecture displays significant improvement over conventional classical and deep algorithms for forensic cross-device speaker verification.

## ACKNOWLEDGEMENT

This work is based upon a work supported by the Center for Identification Technology Research and the National Science Foundation under Grant #1650474.

## References

- [1] Praat software: <http://www.fon.hum.uva.nl/praat/>.
- [2] Voicebox: Speech processing toolbox for matlab: <http://www.ee.ic.ac.uk/hp/staff/dmb/voicebox/voicebox.html>.
- [3] O. Abdel-Hamid, A.-r. Mohamed, H. Jiang, L. Deng, G. Penn, and D. Yu. Convolutional neural networks for speech recognition. *IEEE/ACM Transactions on audio, speech, and language processing*, 22(10):1533–1545, 2014.
- [4] O. Abdel-Hamid, A.-r. Mohamed, H. Jiang, and G. Penn. Applying convolutional neural networks concepts to hybrid nn-hmm model for speech recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4277–4280, 2012.

- [5] C. Bhat, B. Vachhani, and S. K. Kopparapu. Recognition of dysarthric speech using voice parameters for speaker adaptation and multi-taper spectral estimation. In *Proc. Interspeech*, pages 228–232, 2016.
- [6] K. Chen and A. Salman. Extracting speaker-specific information with a regularized siamese deep network. In *Advances in Neural Information Processing Systems*, pages 298–306, 2011.
- [7] Y.-h. Chen, I. Lopez-Moreno, T. N. Sainath, M. Visontai, R. Alvarez, and C. Parada. Locally-connected and convolutional neural networks for small footprint speaker recognition. In *Annual Conference of the International Speech Communication Association*, 2015.
- [8] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2005*, volume 1, pages 539–546, 2005.
- [9] A. Dabouei, H. Kazemi, S. M. Iranmanesh, J. Dawson, and N. M. Nasrabadi. Fingerprint distortion rectification using deep convolutional neural networks. In *International Conference on Biometrics*, 2018.
- [10] 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*.
- [11] L. Deng, O. Abdel-Hamid, and D. Yu. A deep convolutional neural network using heterogeneous pooling for trading acoustic invariance with phonetic confusion. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6669–6673, 2013.
- [12] M. Farrús, A. Garde, P. Ejarque, J. Luque, and J. Hernando. On the fusion of prosody, voice spectrum and face features for multimodal person verification. In *Ninth International Conference on Spoken Language Processing*, 2006.
- [13] M. Farrús, J. Hernando, and P. Ejarque. Jitter and shimmer measurements for speaker recognition. In *Eighth Annual Conference of the International Speech Communication Association*, 2007.
- [14] L. P. Heck, Y. Konig, M. K. Sönmez, and M. Weintraub. Robustness to telephone handset distortion in speaker recognition by discriminative feature design. *Speech Communication*, 31(2-3):181–192, 2000.
- [15] G. Heigold, I. Moreno, S. Bengio, and N. Shazeer. End-to-end text-dependent speaker verification. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5115–5119, 2016.
- [16] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint*, 2015.
- [17] S. M. Iranmanesh, A. Dabouei, H. Kazemi, and N. M. Nasrabadi. Deep cross polarimetric thermal-to-visible face recognition. *arXiv preprint arXiv:1801.01486*, 2018.
- [18] H. Kazemi, S. Soleymani, A. Dabouei, M. Iranmanesh, and N. M. Nasrabadi. Attribute-centered loss for soft-biometrics guided face sketch-photo recognition. In *IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2018.
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [20] A. Larcher, K. A. Lee, B. Ma, and H. Li. Text-dependent speaker verification: Classifiers, databases and rsr2015. *Speech Communication*, 60:56–77, 2014.
- [21] 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*, 2017.
- [22] Y. Liu, T. Fu, Y. Fan, Y. Qian, and K. Yu. Speaker verification with deep features. In *IJCNN*, pages 747–753, 2014.
- [23] M. McLaren, Y. Lei, and L. Ferrer. Advances in deep neural network approaches to speaker recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 4814–4818, 2015.
- [24] S. Motiian, Q. Jones, S. Iranmanesh, and G. Doretto. Few-shot adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 6670–6680, 2017.
- [25] A. Nagrani, J. S. Chung, and A. Zisserman. Voxceleb: a large-scale speaker identification dataset. *arXiv preprint*, 2017.
- [26] H. Nakasone and S. D. Beck. Forensic automatic speaker recognition. In *A Speaker Odyssey-The Speaker Recognition Workshop*, 2001.
- [27] S. J. Park, G. Yeung, J. Kreiman, P. A. Keating, and A. Alwan. Using voice quality features to improve short-utterance, text-independent speaker verification systems. *Proc. Interspeech*, pages 1522–1526, 2017.
- [28] T. N. Sainath, B. Kingsbury, G. Saon, H. Soltau, A.-r. Mohamed, G. Dahl, and B. Ramabhadran. Deep convolutional neural networks for large-scale speech tasks. *Neural Networks*, 64:39–48, 2015.
- [29] D. Snyder, P. Ghahremani, D. Povey, D. Garcia-Romero, Y. Carmiel, and S. Khudanpur. Deep neural network-based speaker embeddings for end-to-end speaker verification. In *IEEE Spoken Language Technology Workshop (SLT)*, pages 165–170, 2016.
- [30] S. Soleymani, A. Dabouei, H. Kazemi, J. Dawson, and N. M. Nasrabadi. Multi-level feature abstraction from convolutional neural networks for multimodal biometric identification. In *24th International Conference on Pattern Recognition (ICPR)*, 2018.
- [31] S. Soleymani, A. Torfi, J. Dawson, and N. M. Nasrabadi. Generalized bilinear deep convolutional neural networks for multimodal biometric identification. In *IEEE International Conference on Image Processing (ICIP)*, 2018.
- [32] E. Variani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez. Deep neural networks for small footprint text-dependent speaker verification. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4052–4056, 2014.
- [33] C. Zhang and K. Koishida. End-to-end text-independent speaker verification with triplet loss on short utterances. In *Proc. of Interspeech*, 2017.