

## Single-channel speech separation with auxiliary speaker embeddings

Shuo Liu, Gil Keren, Björn Schuller

### Angaben zur Veröffentlichung / Publication details:

Liu, Shuo, Gil Keren, and Björn Schuller. 2019. "Single-channel speech separation with auxiliary speaker embeddings." arXiv. Augsburg: Universität Augsburg.  
<https://doi.org/10.48550/arXiv.1906.09997>.



# Single-Channel Speech Separation with Auxiliary Speaker Embeddings

Shuo Liu<sup>1</sup>, Gil Keren<sup>1</sup>, Björn Schuller<sup>1,2</sup>

<sup>1</sup>ZD.B Chair of Embedded Intelligence for Health Care and Wellbeing,  
University of Augsburg, Germany

<sup>2</sup>GLAM – Group on Language, Audio & Music, Imperial College London, UK

shuo.liu@informatik.uni-augsburg.de

## Abstract

We present a novel source separation model to decompose a single-channel speech signal into two speech segments belonging to two different speakers. The proposed model is a neural network based on residual blocks, and uses learnt speaker embeddings created from additional clean context recordings of the two speakers as input to assist in attributing the different time-frequency bins to the two speakers. In experiments, we show that the proposed model yields good performance in the source separation task, and outperforms the state-of-the-art baselines. Specifically, separating speech from the challenging VoxCeleb dataset, the proposed model yields 4.79 dB signal-to-distortion ratio, 8.44 dB signal-to-artifacts ratio and 7.11 dB signal-to-interference ratio.

**Index Terms:** single-channel source separation, speech enhancement, deep learning, residual neural network

## 1. Introduction

In the presence of two overlapping speech sources, the human brain is capable of focusing on a selected target speaker and ignoring speech from the other speaker to a large degree. However, constructing an automatic source separation system to extract a target speech signal from the mixture of target and interference speech signals remains a challenging task. The task becomes even more challenging when the two speakers share similar pronunciation and acoustic features. Conventional signal processing algorithms for this task are broadly categorised into multi-channel and single-channel methods, depending on the topology of the microphones used for the signal recording. Techniques based on a microphone-array such as principle component analysis (PCA), independent component analysis (ICA) and non-negative matrix factorisation (NMF) have been reported in the literature to be effective, but these techniques require some additional assumptions such as source independence, space sparsity and non-negative constraints [1, 2, 3, 4]. On the other hand, traditional single-channel source separation approaches such as computational auditory scene analysis (CASA) exploits pitch and onset as grouping cues to decompose a mixed speech signal [5].

Over the past few years, neural network models have managed to outperform classic signal processing algorithms in many speech related applications, such as automatic speech recognition (ASR) [6, 7], speech enhancement [8, 9, 10], and speech emotion recognition (SER) [11, 12, 13]. This recent success of neural network models has inspired some of the recent solutions for the single-channel source separation task [14, 15, 16]. A prominent recent work introduces the deep clustering approach [17]. In this work, an assumption is made that each time-frequency bin in the speech spectrogram belongs to only one speaker. A deep recurrent network is trained to produce

an embedding vector for each time-frequency bin (TF-bin), and the model is trained to emit similar embeddings for TF-bins that originate from the same speaker. A k-means clustering algorithm is then applied at test time to attribute the TF-bins to the different speakers according to the learnt embeddings. The deep attractor network (DaNet) [18] extends the deep clustering framework. In this work, embeddings for the different TF-bins that originate from the same speaker are averaged at training time to create an attractor for this speaker. The model is trained to produce embeddings for the TF-bins that are similar to their respective attractors. The DaNet was shown in experiments to yield superior source separation performance compared to the deep clustering method.

However, both the deep clustering and DaNet methods are facing a non-trivial challenge. Both models have to learn to produce similar embeddings to TF-bins that belong to similar speakers, without any indication in what region in the embedding space the similar embeddings should be. As embeddings for TF-bins that belong to different speakers may concentrate in different regions of the embedding space, given a new speaker at test time, different regions in the audio segment will have to agree in which area of the embedding space the embeddings for this speaker should concentrate. In order for this to happen, bottom layers of the neural network will have to exchange information between all the different parts of the audio segment, communicating the target area in the embedding space. As the network is not explicitly guided to do so, learning this behaviour from data alone is indeed a non-trivial challenge.

In this paper, we propose a novel source separation model based on stacks of residual blocks. To alleviate the above issue, we supply the model with short additional context recordings containing clean speech from each of the speakers. Each of the context recordings is then processed with a dedicated subnetwork to create a speaker embedding vector. The main subnetwork processes the mixed speech segment and emits the separated speech, and is also conditioned on the speaker embedding vectors, that assist in attributing the different TF-bins to the different speakers, thus circumventing the issue described above. It is important to note that supplying the network with those additional inputs is a realistic setting. Indeed, in many real-life applications the audio segments for separation contain large parts where only a single speaker is present. Those parts can be extracted and used by the model as the additional speaker context. In experiments, we find that our proposed model outperforms the state-of-the-art deep clustering and DaNet models in the source separation task using a large-scale speech dataset.

The rest of the paper is organised as follows. Section 2 presents the data used in this work and its processing. Section 3 discusses the structure of the proposed model, experiments and evaluation for the source separation task are described in Section 4. Final conclusions are drawn in Section 5.

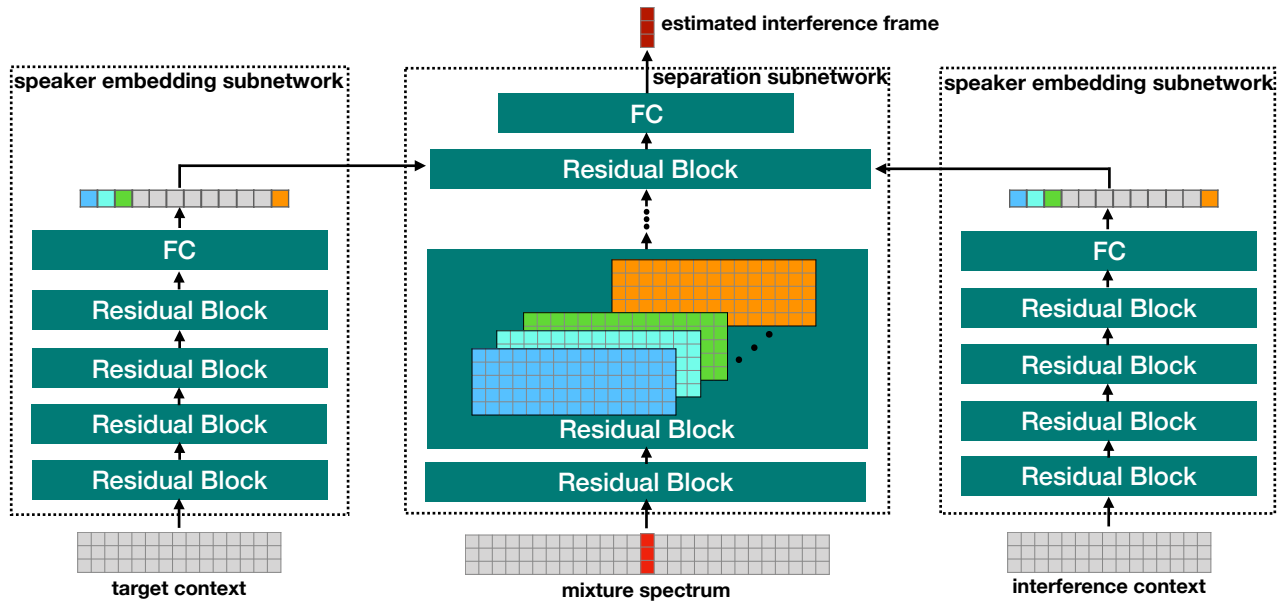


Figure 1: The source separation model architecture. The identical speaker embedding subnetworks process the target and interference contexts via a sequence of 4 residual blocks to produce target and interference speaker embeddings. The separation subnetwork processes the mixture segment through a sequence of 8 residual blocks, each additionally conditioned on the target and interference speaker embeddings, to output an estimated interference frame. The difference between the central frame of the mixture segment and the estimated interference frame is the estimated target frame.

## 2. Data description and processing

In general, the performance of a deep neural network model for source separation improves as the size and diversity of the speech data increases. The VoxCeleb dataset [19, 20] provides more than 2000 hours of single-channel recordings extracted from Youtube interviews of more than 7000 speakers, and includes more than one million utterances. The dataset contains two versions, VoxCeleb1 and VoxCeleb2, each with its own training and test set. As each version consists of distinct speakers, our training and test sets are comprised of the union of the two corresponding sets in the two versions. For the creation of the validation set, 20 utterances from different speakers were chosen from the training set. The speakers with utterances appearing in the validation set were removed from the training set. The above procedure ensures that the training, validation and test sets are speaker independent. The amount of data for validation is selected to be very small because we do not require much hyper-parameters tuning for our proposed source separation network. The utterances in VoxCeleb cover different nationalities and range from 4 to 12 seconds in length.

At training time, the model inputs are created using a random process at each iteration, for improving the separation quality by creating a larger effective size for the training set. At each iteration, we randomly sample two speakers, named the *target speaker* and the *interference speaker*, and a random utterance from each of the two speakers, named the *target utterance* and the *interference utterance* respectively. The two utterances are truncated to have the same length, and then the interference utterance is mixed into the target utterance using a random SNR of either -5, 0, 5, 10, 15 or 25 dB to create the *mixture utterance*. The log magnitude spectrum is extracted from each of the target, interference and mixture utterances by applying a short-time Fourier transform (STFT) on each segment, using a 25 ms Hanning window shifted by 10 ms. Prior to computing

the logarithm, a small value of  $10^{-5}$  is added to the magnitude of the STFT output, to prevent the model from fitting to imperceptible differences in magnitude. The sampling frequency of all audio segments is 16 kHz, therefore each frame consists of 400 samples and its resulting feature vector is comprised of 201 frequencies.

The source separation model processes three inputs. The *mixture segment* consists of 100 successive frames from the mixture spectrum, chosen randomly at each iteration from the mixture utterance. The parts of the target and interference utterance that were used to create the mixture segment are named the *target segment* and the *interference segment* respectively. The *target context* is a 35 frames segment chosen randomly at each iteration from the part of target utterance that does not appear in the mixture segment. The *interference context* is chosen from the interference utterance using the same process. The target and the interference contexts each contain speech from a single speaker, and are used by the model to create the speaker embeddings to assist with the source separation process, as described in Section 3. It is important to note that in many real-life applications the audio segment for separation contains large parts where only a single speaker is present, therefore conditioning our model on the target and interference contexts, that contain speech from a single speaker each, is a realistic setting that may be deployed in a variety of real-life applications. The target segment is used by the separation model as the training label.

For creating the validation and test sets, all utterances from the corresponding set were randomly split into pairs of target and interference utterances, such that each pair contains utterances from different speakers. Each pair was then mixed to create the mixture segment, using an SNR of either -5, -3, -1, 0, 1, 3 or 5 dB. These SNR values ensure a fair comparison of our algorithm with previous work. A larger variety of SNRs were used in the training process, to facilitate the trained model's robustness to a variety of SNRs. Target and interference contexts

were chosen to be from the beginning of the target and interference utterances, respectively. Test and validation sets were created once, and were consistent across all experiments.

### 3. Source separation model

Residual neural networks (resnets) introduce shortcut connections to the conventional CNN framework and enable a substantially deeper architecture, which has been validated to be successful in both the computer vision and audio domains [21, 22, 23]. A basic residual block contains two convolutional layers, where batch normalisation [24] followed by a rectified linear unit (ReLU) are applied between the convolutional layers. The residual block’s input is added to the output of the second convolutional layer, and again batch normalisation and ReLU activation are applied to emit the block’s output. In this work, we use two-dimensional convolutional layers that operate on the time and frequency axes.

The architecture of our proposed source separation model is based on stacks of residual blocks as depicted in Figure 1. The model consists of three subnetworks, each mainly contains a sequence of residual blocks. First, a speaker embedding subnetwork processes the target context, to emit the *target speaker embedding*. Similarly, an identical subnetwork processes the interference context to emit the *interference speaker embedding*. The separation subnetwork then processes the two speaker embeddings and the mixture segment to emit an *estimated interference frame*. The estimated interference frame is an estimate of the interference components in the central frame of the mixture segment. Finally, the *estimated target frame* is computed to be the difference between the central frame of the mixture segment and the estimated interference frame. The model is trained to minimise the squared difference between the estimated target frame and the true target frame that is the central frame of the target segment.

#### 3.1. Speaker embedding subnetwork

Each of the two speaker embedding subnetworks takes a speech context as input, and its output is a speaker embedding vector that may contain valuable acoustic information obtained from the speech segment. The speaker embedding subnetwork is comprised of a sequence of four residual blocks using the specifications from Table 1. The output feature maps of the last residual block is averaged across all locations (time steps and frequency bins) to get a speaker embedding vector with a fixed length of 512.

Table 1: *The speaker embedding subnetwork specifications. The subnetwork contains 4 residual blocks, each with different kernel size, stride, and number of channels.*

| Block | Kernel | Stride | #Channels |
|-------|--------|--------|-----------|
| 1     | (8, 4) | (3, 2) | 64        |
| 2     | (8, 4) | (3, 2) | 128       |
| 3     | (4, 4) | (1, 1) | 256       |
| 4     | (4, 4) | (1, 2) | 512       |

Two identical such subnetworks (each with its own learnable parameters) process the the target and interference contexts and produce the target and interference speaker embeddings. The learnt target and interference speaker embeddings are injected into the separation subnetwork to assist with the source

separation task.

#### 3.2. Separation subnetwork

This separation subnetwork is comprised of a sequence of 8 residual blocks that process the mixture segment, using the specifications from Table 2. For each convolutional layer, the learnt target and interference speaker embeddings are each linearly projected to a dimension equal to the number of feature maps in this layer by applying a trainable fully-connected layer. The two projected vectors are then added to every location in the output map of the convolutional layer. By doing so, all convolutional layers in the separation subnetwork are conditioned on the information of target and interference speech, allowing the separation network to better estimate the components belonging to the interference speaker.

Table 2: *The separation subnetwork specifications. The subnetwork contains 8 residual blocks, each each with different kernel size, stride and number of channels.*

| Block | Kernel | Stride | #Channels |
|-------|--------|--------|-----------|
| 1     | (4, 4) | (1, 1) | 64        |
| 2     | (4, 4) | (1, 1) | 64        |
| 3     | (4, 4) | (2, 2) | 128       |
| 4     | (4, 4) | (1, 1) | 128       |
| 5     | (3, 3) | (2, 2) | 256       |
| 6     | (3, 3) | (1, 1) | 256       |
| 7     | (3, 3) | (2, 2) | 512       |
| 8     | (3, 3) | (1, 1) | 512       |

At last, we flatten the output of the last residual block and feed it through a fully-connected layer to get a 201 dimensional output. This output is added to the central frame of the input mixture segment to yield the estimated target frame. During the training phase, we optimise the network parameters to minimise the mean squared error (MSE) between the estimated target frame and the central frame of the target segment, using stochastic gradient descent (SGD) with a learning rate of 0.1. At validation and test time, inverse Short-Time Fourier transform (iSTFT) was used to reconstruct the target speech using the phase of the mixture segment.

## 4. Experiments and results

We conduct experiments for evaluating the effectiveness of the proposed source separation model. The performance of both our proposed model and the state-of-the-art baselines recently proposed for source separation [17, 18] are compared in a large-scale source separation task using the VoxCeleb dataset and unseen speakers at test time as was described in Section 2.

#### 4.1. Baseline models

As a baseline, we implemented the deep clustering model and the deep attractor network (DaNet) as described in [17, 18]. For experiments with the two baseline methods, we follow the audio preprocessing pipeline from these works. All audio segments were downsampled to 8 kHz before processing and then mixed using one of the SNRs -5, -3, -1, 0, 1, 3, 5 dB. Log magnitude spectrum of the mixture speech was obtained using STFT with a 32 ms window length with an 8 ms window shift. Chunks of 100 frames were used as input to the model. The deep clustering model uses 2 bidirectional LSTM layers, while DaNet




Figure 2: An example of true target and interference spectrum, and the corresponding output of the source separation model.

contains 4 bidirectional LSTM layers, both with 600 hidden units in each layer, to learn 20-dimensional embeddings for every TF-bin of the mixture spectrum. The deep clustering model is trained to produce similar embeddings to TF-bins that originate from the same speaker. The DaNet model averages the embeddings that originate from similar speakers to create attractors, and the model is trained to produce embeddings that are similar to their appropriate attractors. Optimisation parameters were chosen using the validation set. At test time, k-means clustering is used to allocate the TF-bins to the two speakers based on their embedding vectors.

#### 4.2. Comparison

We consider three objective evaluation metrics for separation as described in [25]: signal-to-distortion ratio (SDR), signal-to-artifacts ratio (SAR) and signal-to-interference ratio (SIR). We use the BSSEval toolbox [26] to compute the three evaluation metrics for the VoxCeleb test set and each of the evaluated models. Table 3 contains the results for all the evaluated models. As seen in the results table, our proposed model outperforms the deep clustering and DaNet baselines by a large margin in terms of SDR and SAR, which measures the lack of distortion and algorithm artifacts in the recovered speech signals, concerning speech quality. Regarding SIR, which represents the ability of the model to suppress the interference speech, the DaNet model yielded a better value compared to our approach. However, SIR does not measure whether the target speech was preserved, therefore a high SIR value with a low SAR and SDR value indicates that the DaNet model suppresses both the interference and the target speakers to a large extent. We note that both baseline approaches were reported to yield good source separation performance in previous work. However, in this work we use a challenging dataset of natural speech, which may be the reason for the relatively low performance of these methods in this work.

To better visualise the performance of our source separation algorithm, we depict an example of target and interference

Table 3: Comparison of test set evaluation metrics (in dB). DC and DaNet stand for deep clustering and deep attractor network. SDR, SAR, and SIR represent signal-to-distortion, signal-to-artifacts, and signal-to-interference ratio, respectively.

|                  | SDR         | SAR         | SIR          |
|------------------|-------------|-------------|--------------|
| <b>DC[17]</b>    | 0.84        | 2.09        | 6.58         |
| <b>DaNet[18]</b> | 1.81        | 3.29        | <b>10.41</b> |
| <b>Proposed</b>  | <b>4.79</b> | <b>8.44</b> | 7.11         |

segments from two speakers in the test set on the left column of Figure 2, where the target speaker is male and the interference speaker is female. The right column shows our model’s output, the estimated (separated by the model) target and interference segments.

## 5. Conclusions

In this paper, we developed a single-channel source separation model that uses additional conditioning on separate speaker context recordings. The model learns to create speaker embeddings for unseen speakers from additional context recordings. The speaker embeddings may contain important acoustic information regarding the different speakers, and assist the model in enhancing the source separation quality. In experiments, our proposed model considerably outperforms two state-of-the-art baseline models. Future work should focus on extending this method to operate with any number of speakers, and further improving separation quality on the challenging dataset that was used. Moreover, more work should be done on maintaining speech intelligibility in the source separating process, to overcome distortions that are occasionally introduced in the audio. Furthermore, the speaker embeddings may be further investigated, for better understanding how the model codes the different speakers, and what specific acoustic information assists the model with the separation process.

## 6. References

- [1] S. Winter, H. Sawada, and S. Makino, "Geometrical interpretation of the PCA subspace approach for overdetermined blind source separation," *EURASIP Journal on Advances in Signal Processing*, vol. 2006, no. 1, 2006, 11 pages.
- [2] N. Mitianoudis and M. E. Davies, "Audio source separation: Solutions and problems," *International Journal of Adaptive Control and Signal Processing*, vol. 18, no. 3, pp. 299–314, 2004.
- [3] F. Weninger and B. Schuller, "Optimization and parallelization of monaural source separation algorithms in the openBliSSART toolkit," *Journal of Signal Processing Systems*, vol. 69, no. 3, pp. 267–277, 2012.
- [4] A. Ozerov and C. Févotte, "Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 3, pp. 550–563, 2010.
- [5] G. Hu and D. Wang, "A tandem algorithm for pitch estimation and voiced speech segregation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 8, pp. 2067–2079, 2010.
- [6] Z. Zhang, F. Weninger, M. Wöllmer, J. Han, and B. Schuller, "Towards intoxicated speech recognition," in *Proc. International Joint Conference on Neural Networks (IJCNN)*, Anchorage, Alaska, 2017, pp. 1555–1559.
- [7] M. Seltzer, D. Yu, and Y. Wang, "An investigation of deep neural networks for noise robust speech recognition," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Vancouver, Canada, 2013, pp. 7398–7402.
- [8] G. Keren, J. Han, and B. Schuller, "Scaling speech enhancement in unseen environments with noise embeddings," in *Proc. CHiME Workshop on Speech Processing in Everyday Environments*, Hyderabad, India, 2018, pp. 25–29.
- [9] S. Araki, T. Hayashi, M. Delcroix, M. Fujimoto, K. Takeda, and T. Nakatani, "Exploring multi-channel features for denoising-autoencoder-based speech enhancement," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane, Australia, 2015, pp. 116–120.
- [10] F. Weninger, H. Erdogan, S. Watanabe, E. Vincent, J. Le Roux, J. R. Hershey, and B. Schuller, "Speech enhancement with LSTM recurrent neural networks and its application to noise-robust ASR," in *Proc. International Conference on Latent Variable Analysis and Signal Separation (LVA/ICA)*, Liberec, Czech Republic, 2015, pp. 91–99.
- [11] Z. Zhang, J. Geiger, J. Pohjalainen, A. E.-D. Mousa, W. Jin, and B. Schuller, "Deep learning for environmentally robust speech recognition: An overview of recent developments," *ACM Transactions on Intelligent Systems and Technology*, vol. 9, no. 5, 2018, 14 pages.
- [12] J. Han, Z. Zhang, F. Ringeval, and B. Schuller, "Reconstruction-error-based learning for continuous emotion recognition in speech," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, LA, 2017, pp. 2367–2371.
- [13] P. Tzirakis, G. Trigeorgis, M. A. Nicolaou, B. Schuller, and S. Zafeiriou, "End-to-end multimodal emotion recognition using deep neural networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1301–1309, 2017.
- [14] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, 2018.
- [15] Y.-S. Lee, C.-Y. Wang, S.-F. Wang, J.-C. Wang, and C.-H. Wu, "Fully complex deep neural network for phase-incorporating monaural source separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, LA, 2017, pp. 281–285.
- [16] S. Samui, I. Chakrabarti, and S. K. Ghosh, "Deep recurrent neural network based monaural speech separation using recurrent temporal restricted boltzmann machines," in *Proc. Annual Conference of the International Speech Communication Association (INTER-SPEECH)*, Stockholm, Sweden, 2017, pp. 3622–3626.
- [17] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Shanghai, China, 2016, pp. 31–35.
- [18] Z. Chen, Y. Luo, and N. Mesgarani, "Deep attractor network for single-microphone speaker separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, LA, 2017, pp. 246–250.
- [19] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Proc. Annual Conference of the International Speech Communication Association (INTER-SPEECH)*, Stockholm, Sweden, 2017, pp. 2616–2620.
- [20] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, Hyderabad, India, 2018, pp. 1086–1090.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, 2016, pp. 770–778.
- [22] H. K. Vydana and A. K. Vuppala, "Residual neural networks for speech recognition," in *Proc. European Signal Processing Conference (EUSIPCO)*, Kos island, Greece, 2017, pp. 543–547.
- [23] H. Jung, M.-K. Choi, J. Jung, J.-H. Lee, S. Kwon, and W. Young Jung, "Resnet-based vehicle classification and localization in traffic surveillance systems," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, 2017, pp. 61–67.
- [24] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. International Conference on Machine Learning (ICML)*, Lille, France, 2015, pp. 448–456.
- [25] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [26] F.-R. Stöter, A. Liutkus, and N. Ito, "The 2018 signal separation evaluation campaign," in *Proc. International Conference on Latent Variable Analysis and Signal Separation (LVA/ICA)*, Guildford, UK, 2018, pp. 293–305.