

# A FRAMEWORK FOR ROBUST SPEAKER VERIFICATION IN HIGHLY NOISY ENVIRONMENTS LEVERAGING BOTH NOISY AND ENHANCED AUDIO

*Adam Katav   Yair Moshe   Israel Cohen*

Signal and Image Processing Laboratory (SIPL)  
Andrew and Erna Viterbi Faculty of Electrical & Computer Engineering  
Technion – Israel Institute of Technology, Haifa, Israel  
<https://sipl.ece.technion.ac.il>

## ABSTRACT

Recent advancements in speaker verification techniques show promise, but their performance often deteriorates significantly in challenging acoustic environments. Although speech enhancement methods can improve perceived audio quality, they may unintentionally distort speaker-specific information, which can affect verification accuracy. This problem has become more noticeable with the increasing use of generative deep neural networks (DNNs) for speech enhancement. While these networks can produce intelligible speech even in conditions of very low signal-to-noise ratio (SNR), they may also severely alter distinctive speaker characteristics. To tackle this issue, we propose a novel neural network framework that effectively combines speaker embeddings extracted from both noisy and enhanced speech using a Siamese architecture. This architecture allows us to leverage complementary information from both sources, enhancing the robustness of speaker verification under severe noise conditions. Our framework is lightweight and agnostic to specific speaker verification and speech enhancement techniques, enabling the use of a wide range of state-of-the-art solutions without modification. Experimental results demonstrate the superior performance of our proposed framework.


**Index Terms**— Speaker verification, speaker recognition, speaker embedding, acoustic noise, noise robustness.

## 1. INTRODUCTION

Speaker verification aims to determine whether two audio samples originate from the same speaker. Typical speaker verification applications include voice authentication for personal smart devices, authenticating callers in call centers, securing access in telephone banking, and law enforcement investigations. Speaker verification systems rely on speaker embeddings, which are high-dimensional features representing a speaker's vocal tract characteristics and speaking style. Once created, these embeddings can be compared to determine if two audio samples come from the same person [4, 5].

Over the last decade, the superior feature extraction capabilities of deep neural networks (DNNs) have made them the leading approach for speaker embedding. Among the pioneering works in this area was the introduction of the Time Delay Neural Network (TDNN) [6]. One of the most successful architectures for speaker verification using this approach was developed in [7]. In this design, input mel-frequency coefficients (MFCC) are fed into a TDNN, and the resulting embeddings are termed x-vectors. SpeakerNet [1] is a lightweight model that uses an x-vector-based statistics pooling layer. It comprises residual blocks with 1-dimensional depth-wise separable convolutions, batch normalization, and ReLU layers. Another variant of the x-vector architecture with improved results is the emphasized channel attention, propagation, and aggregation TDNN (ECAPA-TDNN) [8]. This model integrates insights from computer vision, introducing enhancements such as 1D Res2Net modules with skip connections to improve the capture of temporal relationships, Squeeze-and-Excitation layers to emphasize informative channels, refining feature discrimination, and channel attention propagation and aggregation to further distribute attention weights across TDNN layers.

Speech recordings in the real world are often affected by background noise and reverberation. These disturbances make it challenging for a person to understand the speech and can also degrade the quality of extracted speaker embeddings, leading to inaccuracies in speaker verification. Speech enhancement plays an important role in speech signal processing as it aims to improve the intelligibility and quality of speech by removing noise from corrupted signals [9]. While it might seem intuitive that applying speech enhancement to noisy recordings would improve speaker embeddings, this is often untrue. The primary objective of speech enhancement is noise suppression, and it does not explicitly guarantee improvements in downstream tasks such as speaker verification. The artifacts and distortions introduced during the enhancement process can sometimes degrade speaker verification performance even further [10]. This issue has become more prominent recently as generative DNNs are increasingly used



**Fig. 1:** t-SNE visualization of SpeakerNet [1] embeddings extracted from two speakers in the VoxCeleb1 [2] dataset with babble noise at an SNR of -15 dB: (a) without any processing, (b) enhanced using DeepFilterNet3 [3], and (c) generated using the proposed method, which combines embeddings from (a) and (b).

for speech enhancement. These models can produce superior speech quality and effectively enhance speech contaminated by higher noise levels. However, their generative nature can lead to significant distortions of the speaker’s intrinsic characteristics in the speech signal, especially under challenging noise conditions.

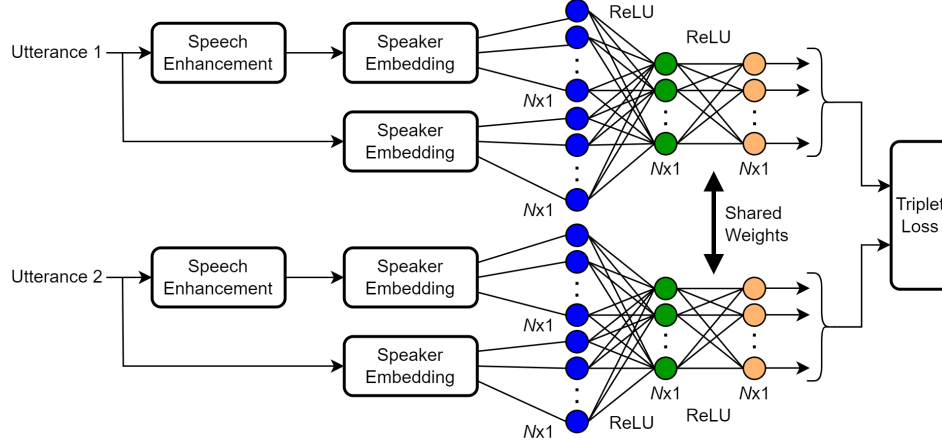
Several previous studies have investigated speaker verification in noisy conditions. Some of these works propose training a speech enhancement module specifically for speaker verification, using a tailored loss function [11] or learning a mapping from noisy to clean speech embeddings [12, 13]. Other approaches focus on learning to extract robust speaker embeddings by separately capturing noise and speaker characteristics [14, 15]. A common strategy is to use a cascaded architecture, where instead of separately processing speech enhancement and speaker verification, the two modules are integrated into a single framework through joint optimization [16, 17, 18]. While existing works demonstrate enhanced noise robustness for speaker verification, they rely on training a dedicated speech enhancement module, a speaker verification module, or both. This is a major drawback as training state-of-the-art enhancement and verification modules often demands significant computational resources and access to large high-quality datasets, which may not be readily available. Our proposed framework offers a more practical solution by utilizing any such pre-trained enhancement or verification module out-of-the-box. This approach not only simplifies the design but also significantly reduces computation complexity. Unlike [19], which employs a learning-based interpolation agent to automatically determine the optimal linear combination of noisy and enhanced signals, our framework combines the speaker embeddings extracted from both sources. This nonlinear combination of embeddings is performed in a highly informative latent space, enabling improved performance without the added complexity of training an interpolation agent. Furthermore, our uti-

lization of state-of-the-art speech enhancement techniques allows us to deliver reliable speaker verification performance even in severe noisy conditions where previous speaker verification methods fail and where speech enhancement methods introduce significant distortions to speaker characteristics.

## 2. ROBUST SPEAKER VERIFICATION

The solution we propose is based on the understanding that speaker embeddings extracted from noisy speech signals and their corresponding enhanced signal embeddings provide complementary information. By combining this complementary information, we can create a more robust embedding that is resistant to noise. In conditions with relatively low noise, the noisy embedding is expected to be more informative because it is less likely to be affected by significant noise artifacts. On the other hand, in challenging noise conditions, the enhanced embedding is likely to provide more valuable information, benefiting from the noise reduction achieved through the enhancement process. We suggest training a neural network to effectively combine these two types of embeddings, taking into account various types and levels of noise. This network can adaptively adjust the contributions of each embedding based on the noise level perceived in the input speech. As a result, this approach can enhance speaker verification performance across different noise scenarios. Figure 1 illustrates this concept with a t-SNE visualization of embeddings from noisy utterances of two speakers in the VoxCeleb1 dataset. The figure shows that, compared to our approach where noisy and enhanced embeddings of each utterance are combined, the noisy and enhanced embeddings of the two speakers are more distinctly separable. This facilitates more accurate speaker verification.

As illustrated in Fig. 2, our framework is based on Siamese neural networks. These neural networks are specifically designed to compare pairs of inputs and assess their



**Fig. 2:** Proposed Siamese architecture for learning robust speaker embeddings.

degree of similarity. This capability has proven invaluable in face recognition, image matching, and fingerprint verification applications. A Siamese architecture involves feeding two inputs into the network: a query input and a reference input. Each input is processed by identical subnetworks with the same weights and parameters, resulting in extracted features. A distance function is then applied to calculate the distance between the feature vectors of the two inputs. The computed distance is a metric for determining the similarity between the inputs, with a smaller distance indicating a higher degree of similarity. In our case, the inputs to the network are two speech utterances. Two embeddings of size  $N \times 1$  are extracted for each speech utterance: one from the noisy utterance and one from the enhanced utterance. The two embeddings of each utterance are combined using a 3-layer MLP (Multi-Layer Perceptron). The first layer is of size  $2N \times 1$ , while the other two are each of size  $N \times 1$ . Between the layers, a ReLU activation function is applied. The MLP learns nonlinear relationships between the two input embeddings and performs dimensionality reduction to output a robust speaker embedding of size  $N \times 1$ . The two MLPs, which follow the Siamese architecture, share the same structure and weights. These weights are optimized based on the similarity between the two robust embeddings, as determined by the loss function. Note that the 3-layer MLP is significantly more straightforward to train than deeper networks and requires less computational resources.

This framework offers the flexibility to employ any speech enhancement and speaker embedding method, enabling the seamless integration of state-of-the-art techniques out-of-the-box. By leveraging cutting-edge speech enhancement methods, we can deliver reliable speaker verification performance even in highly challenging noisy environments, where other speaker verification methods struggle and speech enhancement methods introduce significant distortions to speaker characteristics.

We employ a triplet loss function [20] commonly used for recognition and verification tasks. This loss aims to learn a distance metric that effectively distinguishes between similar and dissimilar examples. We use a variant of triplet loss based on cosine distance to account for the magnitude-invariance property often desired when comparing speaker embeddings. This loss is defined as:

$$\mathcal{L}(A, P, N) = \max(0, d(A, P) - d(A, N) + \alpha) \quad (1)$$

where  $A$  is an anchor utterance,  $P$  is a utterance from the same speaker as  $A$ , and  $N$  is a utterance from a different speaker. The margin  $\alpha$  controls the separation between positive and negative pairs. The function  $d(X, Y)$  is the cosine distance between  $X$  and  $Y$ , derived from the cosine similarity that measures the angle  $\theta$  between the two vectors:

$$d(X, Y) = 1 - \cos(\theta) = 1 - \frac{X \cdot Y}{\|X\| \|Y\|}. \quad (2)$$

Triplet loss minimizes the distance between an anchor and its positive examples (utterances from the same speaker), while maximizing the distance between the anchor and its negative examples (utterances from different speakers). This approach encourages the model to place utterances from the same speaker closer together in the embedding space, while ensuring that utterances from different speakers are positioned further apart.

### 3. RESULTS

Our proposed framework was trained and evaluated using the VoxCeleb1 dataset [2], a collection of celebrity utterances extracted from YouTube videos. The VoxCeleb1 training set comprises 148,642 utterances from 1,211 speakers, while the test set contains 4,874 utterances from 40 speakers. To simulate real-world noise conditions, we augmented both training and test sets with recordings from the MUSAN corpus

**Table 1:** Speaker verification results using our proposed method, compared to verification on noisy signals and noisy signals enhanced with DeepFilterNet3 [3]. The evaluation was performed using SpeakerNet [1] and ECAPA-TDNN [8] speaker embeddings, with noise from the MUSAN corpus [21], covering three noise types at varying SNR levels. Performance is measured using the Equal Error Rate (EER), with the best results highlighted in bold.

		SpeakerNet			ECAPA-TDNN		
Type	SNR	Noisy	Enhc	Ours	Noisy	Enhc	Ours
Noise	0	<b>9.70</b>	13.45	13.17	<b>3.31</b>	7.68	9.43
	-5	16.39	18.18	<b>15.67</b>	<b>5.50</b>	12.22	10.76
	-10	26.19	25.12	<b>19.31</b>	<b>12.53</b>	19.08	14.74
	-15	34.71	32.77	<b>25.21</b>	23.23	26.91	<b>21.48</b>
	-20	41.66	39.69	<b>33.00</b>	31.86	34.27	<b>29.90</b>
Music	0	<b>12.42</b>	14.62	15.33	<b>4.96</b>	8.31	12.19
	-5	23.73	22.80	<b>19.48</b>	<b>13.61</b>	16.28	16.37
	-10	36.82	33.72	<b>27.37</b>	28.41	27.38	<b>24.49</b>
	-15	44.46	41.86	<b>36.37</b>	41.11	38.09	<b>34.96</b>
	-20	48.69	47.39	<b>44.05</b>	47.90	45.77	<b>43.88</b>
Babble	0	<b>20.24</b>	24.79	21.89	<b>19.13</b>	23.16	24.04
	-5	32.64	36.59	<b>26.82</b>	34.89	37.63	<b>31.43</b>
	-10	43.85	44.77	<b>32.30</b>	45.15	45.50	<b>38.47</b>
	-15	46.72	47.48	<b>37.21</b>	48.27	47.78	<b>42.18</b>
	-20	48.31	48.38	<b>41.73</b>	48.77	48.62	<b>45.38</b>

[21], which includes three categories of noises: 6 hours of general noise, such as DTMF tones, thunder, footsteps, paper rustling, and animal noises, 42 hours of music, and 60 hours of speech babble noise. We divided the MUSAN corpus into two disjoint sets: one for training set augmentation and the other for test set augmentation. Each utterance in the training set was corrupted with a randomly selected signal-to-noise ratio (SNR) between 0 and -20 dB. The test set was evaluated at SNR levels of {0, -5, -10, -15, -20} dB. These SNR levels are lower than those typically found in previous literature. Our framework is especially beneficial at these lower SNR levels due to its innovative use of state-of-the-art speech enhancement and speaker verification models out-of-the-box. Our model was trained using the AdamW optimizer with a batch size of 32, a learning rate of  $10^{-3}$ , and a triplet margin parameter  $\alpha = 0.25$ . Since it is lightweight, it was trained for just 10 minutes on an NVIDIA RTX 3090.

Table 1 summarizes the results of our proposed method, compared to speaker verification using noisy and enhanced embeddings, where the enhancement was performed with DeepFilterNet3 [3]. The evaluation was conducted using speaker embeddings from SpeakerNet [1] and ECAPA-TDNN [8]. Notably, ECAPA-TDNN demonstrates greater robustness to noise compared to SpeakerNet, typically resulting in better performance at lower SNR levels. At SNR = 0, the noisy embedding achieves the best verification performance. However, at lower SNRs, the enhanced embedding occasionally outperforms it, while our proposed method consistently delivers the best results in these lower SNRs.

#### 4. CONCLUSIONS

This paper presented a novel neural network framework for robust speaker verification in challenging acoustic environments. The proposed Siamese architecture effectively integrates speaker embeddings from both noisy and enhanced speech, leveraging their complementary information to improve verification performance. By utilizing state-of-the-art, pre-trained speech enhancement and speaker verification models, the framework eliminates the need for task-specific training, making it both practical and computationally efficient. Experimental results demonstrate the superiority of the proposed approach across diverse noise conditions, particularly in highly degraded environments where conventional speaker verification methods often fail. Moreover, the framework is inherently flexible and agnostic to specific speech enhancement and speaker verification techniques, enabling seamless integration with future advancements in these fields.

#### 5. ACKNOWLEDGMENT

The authors thank Ram Binshtock, Sahar Zeltzer, and David Portal for their contributions in the early stages of this work as part of the Signal Processing Cup 2024 Challenge. This research was supported by the Israel Science Foundation (grant no. 1449/23) and the Pazy Research Foundation.

## 6. REFERENCES

- [1] Nithin Rao Koluguri, Jason Li, Vitaly Lavrukhin, and Boris Ginsburg, “SpeakerNet: 1D depth-wise separable convolutional network for text-independent speaker recognition and verification,” *arXiv preprint arXiv:2010.12653*, 2020.
- [2] Arsha Nagrani, Joon Son Chung, Weidi Xie, and Andrew Senior, “VoxCeleb: Large-scale speaker verification in the wild,” *Computer Speech & Language*, vol. 60, pp. 101027, 2020.
- [3] Hendrik Schröter, Tobias Rosenkranz, Andreas Maier, et al., “DeepFilterNet: Perceptually motivated real-time speech enhancement,” *arXiv preprint arXiv:2305.08227*, 2023.
- [4] Zhongxin Bai and Xiao-Lei Zhang, “Speaker recognition based on deep learning: An overview,” *Neural Networks*, vol. 140, pp. 65–99, 2021.
- [5] Maros Jakubec, Roman Jarina, Eva Lieskovska, and Peter Kasak, “Deep speaker embeddings for speaker verification: Review and experimental comparison,” *Engineering Applications of Artificial Intelligence*, vol. 127, pp. 107232, 2024.
- [6] Kevin J Lang, Alex H Waibel, and Geoffrey E Hinton, “A time-delay neural network architecture for isolated word recognition,” *Neural networks*, vol. 3, no. 1, pp. 23–43, 1990.
- [7] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, “X-vectors: Robust dnn embeddings for speaker recognition,” in *Proc. IEEE Internat. Conf. Acoust. Speech Signal Process., ICASSP*, 2018, pp. 5329–5333.
- [8] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuyne, “ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” *arXiv preprint arXiv:2005.07143*, 2020.
- [9] Asri Rizki Yuliani, M Faizal Amri, Endang Suryawati, Ade Ramdan, and Hilman Ferdinandus Pardede, “Speech enhancement using deep learning methods: A review,” *Jurnal Elektronika dan Telekomunikasi*, vol. 21, no. 1, pp. 19–26, 2021.
- [10] Seyed Omid Sadjadi and John HL Hansen, “Assessment of single-channel speech enhancement techniques for speaker identification under mismatched conditions,” in *Proc. INTERSPEECH*, 2010, pp. 2138–2141.
- [11] Saurabh Kataria, Phani Sankar Nidadavolu, Jesús Vilalba, Nanxin Chen, Paola Garcia-Perera, and Najim Dehak, “Feature enhancement with deep feature losses for speaker verification,” in *Proc. IEEE Internat. Conf. Acoust. Speech Signal Process., ICASSP*, 2020, pp. 7584–7588.
- [12] Oldrich Plchot, Lukas Burget, Hagai Aronowitz, and Pavel Matejka, “Audio enhancing with dnn autoencoder for speaker recognition,” in *Proc. IEEE Internat. Conf. Acoust. Speech Signal Process., ICASSP*, 2016, pp. 5090–5094.
- [13] Danwei Cai, Weicheng Cai, and Ming Li, “Within-sample variability-invariant loss for robust speaker recognition under noisy environments,” in *Proc. IEEE Internat. Conf. Acoust. Speech Signal Process., ICASSP*, 2020, pp. 6469–6473.
- [14] Ya-Qi Yu, Siqi Zheng, Hongbin Suo, Yun Lei, and Wu-Jun Li, “CAM: Context-aware masking for robust speaker verification,” in *Proc. IEEE Internat. Conf. Acoust. Speech Signal Process., ICASSP*, 2021, pp. 6703–6707.
- [15] Yao Sun, Hanyi Zhang, Longbiao Wang, Kong Aik Lee, Meng Liu, and Jianwu Dang, “Noise-disentanglement metric learning for robust speaker verification,” in *Proc. IEEE Internat. Conf. Acoust. Speech Signal Process., ICASSP*, 2023, pp. 1–5.
- [16] Suwon Shon, Hao Tang, and James Glass, “VoiceID loss: Speech enhancement for speaker verification,” *arXiv preprint arXiv:1904.03601*, 2019.
- [17] Ju-ho Kim, Jungwoo Heo, Hye-jin Shim, and Ha-Jin Yu, “Extended U-Net for speaker verification in noisy environments,” *arXiv preprint arXiv:2206.13044*, 2022.
- [18] Yanfeng Wu, Taihao Li, Junan Zhao, Qirui Wang, and Jing Xu, “A fused speech enhancement framework for robust speaker verification,” *IEEE Signal Processing Letters*, 2023.
- [19] Chi-Chang Lee, Hong-Wei Chen, Chu-Song Chen, Hsin-Min Wang, Tsung-Te Liu, and Yu Tsao, “LC4SV: A denoising framework learning to compensate for unseen speaker verification models,” in *Proc. 2023 IEEE Automatic Speech Recognition and Understanding Workshop, ASRU*, 2023, pp. 1–8.
- [20] Florian Schroff, Dmitry Kalenichenko, and James Philbin, “Facenet: A unified embedding for face recognition and clustering,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2015, pp. 815–823.
- [21] David Snyder, Guoguo Chen, and Daniel Povey, “MUSAN: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.