



HAL
open science

Assessing the Impact of Speaker Identity in Speech Spoofing Detection

Anh-Tuan Dao, Driss Matrouf, Nicholas Evans

► **To cite this version:**

Anh-Tuan Dao, Driss Matrouf, Nicholas Evans. Assessing the Impact of Speaker Identity in Speech Spoofing Detection. 2026. <hal-05591353>

HAL Id: hal-05591353

<https://hal.science/hal-05591353v1>

Preprint submitted on 21 Apr 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

ASSESSING THE IMPACT OF SPEAKER IDENTITY IN SPEECH SPOOFING DETECTION

Anh-Tuan DAO, Driss Matrouf

Nicholas Evans

Laboratoire d’informatique d’Avignon, France

EURECOM, Sophia Antipolis, France

ABSTRACT

Spoofing detection systems are typically trained using diverse recordings from multiple speakers, often assuming that the resulting embeddings are independent of speaker identity. However, this assumption remains unverified. In this paper, we investigate the impact of speaker information on spoofing detection systems. We propose two approaches within our *Speaker-Invariant Multi-Task* framework, one that models speaker identity within the embeddings and another that removes it. SInMT integrates multi-task learning for joint speaker recognition and spoofing detection, incorporating a gradient reversal layer. Evaluated using four datasets, our speaker-invariant model reduces the average equal error rate by 17% compared to the baseline, with up to 48% reduction for the most challenging attacks (e.g., A11).

Index Terms— audio deepfake detection, speaker-aware, speaker-invariant, SSL model

1. INTRODUCTION

Automatic Speaker Verification (ASV) systems offer secure biometric recognition based on voice characteristics, but sophisticated spoofing techniques mimicking legitimate users pose significant security challenges. To counter these threats, spoofing detection systems are increasingly deployed to protect ASV integrity. Prior work has explored multi-task learning (MTL) to enhance spoofing detection, with ResNet-based models demonstrating improved performance for ASVspoof 2019 by incorporating speaker recognition as an auxiliary task [1]. These results indicate that speaker-related information provides complementary cues that help discriminate between bona fide and spoofed utterances.

MTL approaches combining speaker recognition and spoofing detection have been more commonly applied in spoofing-aware speaker verification (SASV) tasks [2], but attract less attention in standalone spoofing detection. One reason could be the limited number of speakers in spoofed datasets like ASVspoof 2019 [3] which, despite providing speaker labels, lacks sufficient diversity for robust speaker-related training. Prior work [4] addressed this through multi-stage training: first pre-training models using large-scale speaker recognition datasets like VoxCeleb2, followed by fine-tuning using the ASVspoof 2019 dataset with multi-task

objectives. The recent ASVspoof 5 [5] database, with its expanded set of speaker IDs, supports direct MTL for spoofing detection, motivating our investigation into unified training paradigms optimized jointly for both tasks.

Recent self-supervised learning (SSL) models significantly outperform traditional architectures by leveraging large-scale pre-training to capture rich speech representations [6, 7, 8]. Spoofing detection systems are generally trained on varied recordings from multiple speakers, typically assuming that generated embeddings are independent of speaker identity. This hypothesis is questionable and is dependent on the training data, particularly when the number of speakers or their diversity is limited.

In this paper, we investigate the role of speaker identity in spoofing detection systems. Specifically, our goal is to determine whether explicitly incorporating speaker-related information enhances performance or, conversely, degrades it. Additionally, we explore whether suppressing speaker-specific information could be beneficial. To address these questions, we propose the **Speaker-Invariant Multi-Task (SInMT)** framework, an SSL-based architecture which exploits SSL-based feature extraction with Multi-Head Factorized Attention (MHFA) classifiers. SInMT supports flexible switching between two configurations. The first, **Speaker-Invariant (MHFA-IVspk)**, employs a gradient reversal layer (GRL) with two classification heads, one for binary spoof detection and one for speaker identification (using speaker IDs), to adversarially minimize speaker discriminability. By suppressing speaker-specific information, the model focuses on spoofing indicators. The second, **Speaker-Aware (MHFA-spk)**, jointly trains spoofing detection and speaker identification tasks without GRL, leveraging speaker-dependent cues to improve detection in scenarios where spoofs exhibit mismatched speaker characteristics. The framework enables seamless switching between configurations by activating or deactivating the GRL, allowing direct comparisons of speaker-aware and speaker-invariant strategies within a unified SSL backbone.

Our main contributions are:

- **SInMT framework:** We investigate the impact of speaker information on spoofing detection systems by proposing a flexible SSL-based framework. The SInMT unifies speaker-aware and speaker-invariant

approaches for spoofing detection, built using SSL representations with dual MHFA classifiers. We will show that handling speaker information, either by removing it or integrating it, improves performance.

- **Comprehensive empirical evaluation:** We demonstrate that both configurations consistently outperform the baseline MHFA model, with MHFA-IVspk reducing the average equal error rate (EER) by 17.2% across four datasets, and delivering substantial improvements for the most challenging attack types (e.g., a 48% EER reduction for attack A11).

2. SSL-BASED SPOOFING DETECTION

Our approach to spoofing detection uses SSL models to extract representations from raw audio. The model architecture includes two primary components: an SSL-based feature extractor and a back-end classifier. The feature extraction module comprises a CNN encoder and a Transformer-based contextual network. The CNN encoder transforms raw audio into latent representations $z_{1:T}$, compressing dimensionality. These representations are then processed by the Transformer-based contextual network, yielding context-aware representations $o_{1:T}$. The Transformer’s self-attention mechanism effectively models long-range dependencies critical for interpreting speech patterns. We assess and benchmark three back-end classifiers as baselines: Conformer [7], MHFA [8], and AASIST [9].

3. SPEAKER-INVARIANT MULTI-TASK FRAMEWORK FOR SPOOFING DETECTION

We propose the SInMT framework for speech spoofing detection, designed to either leverage or suppress speaker information through a flexible architecture. SInMT enables two configurations: a speaker-aware (MHFA-spk) that jointly optimizes spoofing detection and speaker recognition, and a speaker-invariant (MHFA-IVspk) that uses a GRL to suppress speaker-specific patterns.

3.1. Architecture

The SInMT framework uses a SSL backbone and MHFA classifiers, as depicted in Figure 1. The MHFA classifiers used in the SInMT framework are identical to those of the MHFA baseline model. SInMT comprises:

- **Feature Extractor:** A pre-trained XLSR encoder [10] that maps raw audio to a sequence of contextualized frame-level embeddings.
- **Spoofing Classifier Head:** A MHFA network that uses the XLSR embeddings to predicts the binary label $\hat{y}_s \in \{\text{bonafide}, \text{spoof}\}$.




Fig. 1: Our model architecture with dual classifiers and GRL.

- **Speaker Classifier Head:** A second MHFA network (structurally identical to the spoofing head) that predicts the speaker ID $\hat{y}_d \in 1, \dots, D$, where D is the number of speakers used during training.
- **Gradient Reversal Layer (GRL):** Inserted between the feature extractor and the speaker classifier head. In the forward pass, the GRL acts as an identity. During backpropagation, it multiplies the gradients by a hyperparameter $-\lambda$ to encourage the feature extractor to remove speaker-specific patterns.

For MHFA-spk, the spoofing and speaker heads are trained jointly with a gradient scaling factor $\lambda = -1$ to leverage speaker-specific cues.

3.2. Speaker-Invariant Adversarial Training

To generate embeddings that are invariant to speaker identity, we employ an adversarial training objective. For an input audio signal x , the feature sequence is denoted as $o = f(x)$. The spoofing classifier is represented by $g_s(\cdot)$, and the speaker classifier by $g_d(\cdot)$. The training process optimizes the following objective:

$$\min_{f, g_s, g_d} \mathcal{L}_s(g_s(f(x)), y_s) + \alpha \mathcal{L}_d(g_d(\text{GRL}(f(x))), y_d) \quad (1)$$

Here, the terms are defined as:

- \mathcal{L}_s : cross-entropy loss for classifying spoofing attempts.
- \mathcal{L}_d : cross-entropy loss for speaker identification.
- α : a hyperparameter that adjusts the balance between spoofing detection and speaker invariance.

3.3. Parameter Optimization via Backpropagation

Model parameters are updated using stochastic gradient descent to achieve speaker-invariant features for spoofing detection. Let θ_f , θ_{g_s} , and θ_{g_d} denote the parameters of the XLSR

Table 1: Model performance comparison in term of Equal Error Rate (EER).

Model	EER				
	ITW	ASV5 eval	ASV21LA	ASV21DF	Pooled
AASIST	7.03	5.54	13.66	9.60	8.95
Conformer	5.69	3.85	12.49	10.40	8.10
MHFA	4.31	4.64	12.14	8.58	7.41
MHFA-spk	3.76	5.29	8.67	8.41	6.53
MHFA-IVspk	3.58	4.98	8.41	7.57	6.13

Table 2: Model performance (% EER) of different models, evaluated on ASVspoof 2021 LA break down: attack types (A07-19). The *pooled* indicates the average performance across all categories.

Model	A07	A08	A09	A10	A11	A12	A13	A14	A15	A16	A17	A18	A19	Pooled
MHFA	1.54	1.91	0.76	20.77	17.02	3.45	4.75	3.82	1.49	4.32	7.01	19.67	37.57	12.14
MHFA-spk	1.07	1.14	0.67	14.40	10.30	3.72	6.61	1.94	1.15	3.72	4.93	11.18	28.61	8.67
MHFA-IVspk	0.88	1.15	0.45	12.56	8.76	3.52	3.61	2.36	1.12	3.32	4.76	15.45	24.85	8.41

feature extractor, spoofing classifier, and speaker classifier, respectively. The updates are performed as follows:

$$\theta_f \leftarrow \theta_f - \mu \left(\frac{\partial \mathcal{L}_s^i}{\partial \theta_f} - \lambda \frac{\partial \mathcal{L}_d^i}{\partial \theta_f} \right) \quad (2)$$

$$\theta_{g_s} \leftarrow \theta_{g_s} - \mu \frac{\partial \mathcal{L}_s^i}{\partial \theta_{g_s}} \quad (3)$$

$$\theta_{g_d} \leftarrow \theta_{g_d} - \mu \frac{\partial \mathcal{L}_d^i}{\partial \theta_{g_d}} \quad (4)$$

In Equations (2) to (4), μ represents the learning rate, and λ is a hyperparameter that controls the weight of the speaker-invariant loss [11]. The feature extractor update in (2) minimizes the spoofing classification loss \mathcal{L}_s^i while maximizing the speaker classification loss \mathcal{L}_d^i through the GRL, which scales the speaker classifier gradient by $-\lambda$ during backpropagation. This aims to confuse the speaker classifier, thereby promoting speaker invariance. Equations (3) and (4) update the spoofing and speaker classifiers, respectively, using standard gradient descent to optimize $g_s(\cdot)$ and $g_d(\cdot)$.

4. EXPERIMENTAL SETUP

4.1. Datasets and Metrics

We train models using the ASVspoof 5 training set which comprises approximately 180,000 utterances collected from 400 speakers. For evaluation, we use four datasets: ASVspoof 5 evaluation set, In-the-Wild (ITW) [12], ASVspoof 2021 LA and DF hidden subsets [13, 14]. These hidden subsets have non-speech segments removed, preventing models from exploiting such shortcuts. Prior work [15] shows the challenges presented by these hidden subsets, compared to standard evaluation subsets, since they demand reliance on core spoofing cues rather than non-speech-related shortcuts. Performance is evaluated using the EER.

4.2. Data Augmentation

During training, we employed common data augmentation strategies using the MUSAN dataset and the real room impulse response (RIR) database [16, 17]. Each training utterance is enhanced with one of four augmentation techniques: reverberation, speech, music, or noise ([18], Section 3.2). For consistency, all experiments reported in this paper adopt the same data augmentation protocol.

4.3. Implementation Details

During training, input utterances are randomly segmented into 4-second clips. For evaluation, models operate on complete audio clips. We use the Adam optimizer [19] with a learning rate of 10^{-6} to minimize a weighted cross-entropy loss. Training is conducted in batches of 32 samples over 30 epochs using NVIDIA A100 GPUs. The hyperparameter α is set to 0.1. We train the MHFA-spk model with a gradient scaling factor $\lambda = -1$, leveraging speaker-specific cues. Subsequently, the MHFA-IVspk model is initialized from the pre-trained MHFA-spk model, and then trained with the GRL-enabled adversarial objective using $\lambda = 1$ to enforce speaker invariance.

5. RESULTS

We compare the performance of three baseline models (AASIST, Conformer, and MHFA), then choose the best (MHFA) for experiments with our SInMT framework.

5.1. Model Performance Comparison

Table 1 presents performance for four datasets. Among the baseline models, MHFA achieves the best average performance, reducing the pooled EER by 17.2% compared to AASIST. The most significant improvement is observed for

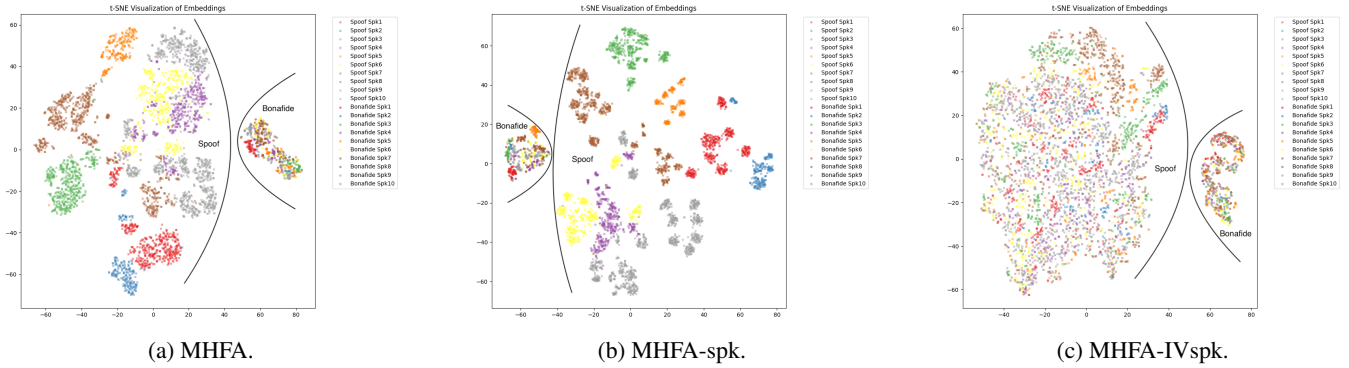


Fig. 2: t-SNE visualization of embeddings from MHFA, MHFA-spk and MHFA-IVspk models, showcasing audio representations of ten distinct speakers from the ASVspooof 5 dataset, each represented by a unique color.

the ITW dataset, where MHFA reduces the EER from 7.03% to 4.31% (38.7% relative gain). Consequently, we apply our SInMT framework to the MHFA classifier to further enhance performance.

We now compare MHFA-spk and MHFA-IVspk against the MHFA baseline. MHFA-spk improves the pooled EER by 11.8% (from 7.41% to 6.53%), with the largest gain for the ASVspooof 2021 LA dataset, reducing the EER from 12.14% to 8.67%. MHFA-IVspk further improves performance, achieving a 17.2% reduction in pooled EER compared to MHFA (from 7.41% to 6.13%). Notable improvements are seen for out-of-domain datasets: ASVspooof 2021 LA (30.7% relative gain, 12.14% to 8.41%) and ITW (17.0% relative gain, 4.31% to 3.58%). MHFA-IVspk slightly outperforms MHFA-spk, but their condition-specific advantages are unclear. We suspect that MHFA-IVspk excels on spoofing datasets with unseen speakers due to its speaker-invariant training, while MHFA-spk may be superior on datasets with seen speakers. As the evaluation datasets feature disjoint speakers, future work will evaluate MHFA-spk on closed-set speaker datasets to confirm these distinctions.

5.2. Visualisation

A visualisation of embeddings using t-SNE is shown in Fig. 2. Specifically, we apply t-SNE on the embeddings of audio samples from 10 distinct speakers selected at random from the ASVspooof 5 training dataset. In the embedding space for the MHFA model, we observe a partial separation between speaker IDs, though clusters remain closely packed. This suggests some retention of speaker-specific information. For the MHFA-spk model, the visualization shows more pronounced clustering by speaker identity, indicating that this model better preserves speaker-specific information. Notably, each speaker ID forms small sub-clusters, potentially reflecting individual speaker characteristics. In contrast, MHFA-IVspk, trained with a GRL to suppress speaker discriminability, shows no distinct speaker clusters, corresponding to the ob-

jective of minimising speaker-specific information while focusing on spoofing-related information cues.

5.3. Attack Types Breakdown Analysis

An analysis of model performance across different attack types in the ASVspooof 2021 LA dataset is presented in Table 2. The dataset includes diverse attack types, ranging from text-to-speech and voice conversion to hybrid approaches. The attack types exhibit performance variability, reflecting their distinct synthesis characteristics. Our proposed models, MHFA-spk and MHFA-IVspk, demonstrate performance gains across all attack types compared to the MHFA baseline. The most substantial improvements are observed for A10 and A11, where the high naturalness and speaker similarity typically pose significant detection challenges. Specifically, compared to MHFA, MHFA-IVspk achieves a 40% relative EER reduction for A10 (from 20.7% to 12.5%) and a 48% relative EER reduction on A11 (from 17.0% to 8.7%).

6. CONCLUSION

This study demonstrates the critical role of speaker-specific information in enhancing spoofing detection systems. We found that handling speaker information, either by removing it or integrating it, is crucial. Our proposed SInMT framework, leveraging the MTL with a GRL, effectively integrates speaker identification and spoofing detection. The SInMT enables flexible switching between speaker-aware (MHFA-spk) and speaker-invariant (MHFA-IVspk) configurations. Assessed across four datasets, both configurations outperform the MHFA baseline, with the MHFA-IVspk model reducing the pooled EER by 17%. Future work will explore a hybrid approach that dynamically integrates speaker-aware and speaker-invariant models.

7. ACKNOWLEDGEMENTS

This project was provided with computing HPC and storage resources by GENCI at IDRIS to thank the grant 2025-AD011016488 on the supercomputer Jean Zay’s A100 partition. This work was financially supported by ANR BRUEL (ANR-22-CE39-0009).

8. REFERENCES

- [1] Yichuan Mo and Shilin Wang, “Multi-task learning improves synthetic speech detection,” in *ICASSP 2022*, 2022.
- [2] Jee-weon Jung, Hemlata Tak, Hye-jin Shim, Hee-Soo Heo, Bong-Jin Lee, Soo-Whan Chung, Ha-Jin Yu, Nicholas Evans, and Tomi Kinnunen, “SASV 2022: The first spoofing-aware speaker verification challenge,” in *Interspeech 2022*, 2022.
- [3] Massimiliano Todisco, Xin Wang, Ville Vestman, Md. Sahidullah, Héctor Delgado, Andreas Nautsch, Junichi Yamagishi, Nicholas Evans, Tomi H. Kinnunen, and Kong Aik Lee, “ASVspoof 2019: Future horizons in spoofed and fake audio detection,” in *Interspeech 2019*, 2019.
- [4] Sung Hwan Mun, Hye-jin Shim, Hemlata Tak, Xin Wang, Xuechen Liu, Md. Sahidullah, Myeonghun Jeong, Min Hyun Han, Massimiliano Todisco, Kong Aik Lee, Junichi Yamagishi, Nicholas Evans, Tomi Kinnunen, Nam Soo Kim, and Jee-weon Jung, “Towards single integrated spoofing-aware speaker verification embeddings,” in *Interspeech 2023*, 2023.
- [5] Xin Wang, Héctor Delgado, Hemlata Tak, Jee weon Jung, Hye jin Shim, Massimiliano Todisco, Ivan Kukanov, Xuechen Liu, Md Sahidullah, Tomi H. Kinnunen, Nicholas Evans, Kong Aik Lee, and Junichi Yamagishi, “ASVspoof 5: crowdsourced speech data, deepfakes, and adversarial attacks at scale,” in *ASVspoof 2024*, 2024.
- [6] Jee-weon Jung, Hee-Soo Heo, Hemlata Tak, Hye-jin Shim, Joon Son Chung, Bong-Jin Lee, Ha-Jin Yu, and Nicholas Evans, “AASIST: audio anti-spoofing using integrated spectro-temporal graph attention networks,” in *ICASSP 2022*. 2022, pp. 6367–6371, IEEE.
- [7] Eros Rosello, Alejandro Gómez Alanís, Angel M. Gomez, and Antonio M. Peinado, “A conformer-based classifier for variable-length utterance processing in anti-spoofing,” in *Interspeech 2023*, 2023.
- [8] Theophile Stourbe, Victor Miara, Théo Lepage, and Réda Dehak, “Exploring WavLM back-ends for speech spoofing and deepfake detection,” in *ASVspoof 2024*, 2024.
- [9] Tak Hemlata, Todisco Massimiliano, Wang Xin, Jung Jee-weon, Yamagishi Junichi, and Nicholas Evans, “Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation,” in *Odyssey 2022*, 2022.
- [10] Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, Alexei Baevski, Alexis Conneau, and Michael Auli, “XLS-R: self-supervised cross-lingual speech representation learning at scale,” in *Interspeech 2022*, 2022.
- [11] Yaroslav Ganin and Victor Lempitsky, “Unsupervised domain adaptation by backpropagation,” in *ICML*, 2015.
- [12] Nicolas M. Müller, Pavel Czempin, Franziska Dieckmann, Adam Froghyar, and Konstantin Böttinger, “Does audio deepfake detection generalize?,” in *Interspeech 2022*. 2022, pp. 2783–2787, ISCA.
- [13] Xuechen Liu, Xin Wang, Md. Sahidullah, Jose Patino, Héctor Delgado, Tomi Kinnunen, Massimiliano Todisco, Junichi Yamagishi, Nicholas Evans, Andreas Nautsch, and Kong Aik Lee, “ASVspoof 2021: Towards spoofed and deepfake speech detection in the wild,” *IEEE ACM Trans. Audio Speech Lang. Process.*, 2023.
- [14] Junichi Yamagishi, Xin Wang, Massimiliano Todisco, Md Sahidullah, Jose Patino, Andreas Nautsch, Xuechen Liu, Kong Aik Lee, Tomi Kinnunen, Nicholas Evans, and Héctor Delgado, “ASVspoof 2021: accelerating progress in spoofed and deepfake speech detection,” in *ASVspoof 2021*, 2021.
- [15] Juan M. Martín-Doñas, Aitor Álvarez, Eros Rosello, Angel M. Gomez, and Antonio M. Peinado, “Exploring Self-supervised Embeddings and Synthetic Data Augmentation for Robust Audio Deepfake Detection,” in *Interspeech 2024*, 2024.
- [16] David Snyder, Guoguo Chen, and Daniel Povey, “MUSAN: A music, speech, and noise corpus,” *ArXiv*, 2015.
- [17] Tom Ko, Vijayaditya Peddinti, Daniel Povey, Michael L. Seltzer, and Sanjeev Khudanpur, “A study on data augmentation of reverberant speech for robust speech recognition,” in *ICASSP 2017*, 2017.
- [18] Anh-Tuan Dao, Mickael Rouvier, and Driss Matrouf, “ASVspoof 5 challenge: advanced resnet architectures for robust voice spoofing detection,” in *ASVspoof 2024*, 2024, pp. 163–169.

- [19] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” in *ICLR*, Yoshua Bengio and Yann LeCun, Eds., 2015.