

Voice Morphing: Two Identities in One Voice

Sushanta K. Pani,¹ Anurag Chowdhury, Morgan Sandler, Arun Ross

Abstract: In a biometric system, each biometric sample or template is typically associated with a single identity. However, recent research has demonstrated the possibility of generating “morph” biometric samples that can successfully match more than a single identity. Morph attacks are now recognized as a potential security threat to biometric systems. However, most morph attacks have been studied on biometric modalities operating in the image domain, such as face, fingerprint, and iris. In this preliminary work, we introduce Voice Identity Morphing (VIM) - a voice-based morph attack that can synthesize speech samples that impersonate the voice characteristics of a pair of individuals. Our experiments evaluate the vulnerabilities of two popular speaker recognition systems, ECAPA-TDNN and x-vector, to VIM, with a success rate (MMPMR) of over 80% at a false match rate of 1% on the Librispeech dataset.

Keywords: Identity Morphing, Morph Attack, Speaker Recognition, Speech Synthesis

1 Introduction

Biometric systems use physical or behavioral traits to recognize individuals [JFR07]. A biometric system acquires a biometric sample of an individual (e.g., voice) using a sensor (e.g., microphone) and extracts a salient feature set (or template). This template is then used to recognize the individual. Typically, a template is associated with a single identity. However, over the past decade, several adversarial techniques, called *morph attacks*, have been developed to create synthetic biometric samples that can successfully match multiple identities [FFM14].² Furthermore, in recent times, DeepFake based synthetic image generators have been used to launch morph attacks on image-based biometric systems, viz., face, fingerprint, and iris, with high success rates [Ve21]. The success of such attacks can potentially lead to compromise of security in sensitive applications where a single biometric ID card could be shared by two or more individuals for nefarious purposes.

Existing literature on morph attacks demonstrates its potency against biometric modalities such as face, fingerprint, and iris [Ve21], [Sc17], [SR21]. For example, landmark-based [MFFM19, RRB16] and deep learning-based [Zh21, Da18] face morph attacks have been shown to be effective against face recognition systems. Similarly, researchers have shown the possibility of launching a morph attack against iris matchers both at the image level [FFM14, SR21] and feature level [FCM16, RB17].

¹ All authors are affiliated with Michigan State University, USA, Corresponding author: Arun Ross (rossarun@cse.msu.edu)

² A related vulnerability known as MasterPrint attack [RMR17] or MasterFace attack [Ng22] has also been studied.

The voice modality, on the other hand, has seemingly been spared from morph attacks until now. The use of voice biometrics is especially relevant in some commercial applications, such as digital voice assistants [Ho18] and telephone banking [MSI01]. The voice morphing attack may be particularly harmful in scenarios where verification of a single identity is essential to proceed. For instance, consider an online spoken language test. In this context, the test-taking system might require the candidate to enroll their voice beforehand to ensure that the same individual appears for the test. This step is typically achieved using a speaker recognition system, designed to prevent an accomplice from taking the test on behalf of the candidate. However, with a voice morphing attack method, the candidate could enroll a morphed combination of their voice and that of an accomplice. This blend would match both identities, allowing the accomplice to take the test on the candidate's behalf by successfully matching their voice to the enrolled morphed template. This situation, coupled with the rapid adoption of voice biometric-enabled devices and services, has heightened interest in understanding their vulnerabilities to morphing attacks. Therefore, it is essential to investigate the viability and success rate of such attacks on popular speaker recognition systems.

In this paper, we propose a voice morphing technique called Voice Identity Morphing (VIM)³ that can synthesize artificial voice samples containing the voice characteristics of a pair of identities. Experimentally we show that the morph voice samples generated from two identities can successfully match target audio samples of both constituent identities using two different popular speaker recognition systems. The proposed method uses the DeepTalk network [CRD21] to extract speaker embeddings from two source identities. Then, it performs a feature-level fusion of the two embeddings producing a new embedding corresponding to the morphed identity. Finally, the morphed embedding is input to a Tacotron 2-based Text-to-Speech synthesizer to generate a morphed audio sample.

The main contributions of this preliminary work are as follows: (a) We propose a voice identity morphing technique capable of generating speech samples that can successfully match two identities within the framework of a speaker recognition system. (b) We evaluate and demonstrate the vulnerability of two popular speaker recognition systems, namely x-vector [Sn18] and ECAPA-TDNN [DTD], to our proposed method. (c) We perform an ablation study to better understand this vulnerability, and we initiate a discussion on potential forensic measures that may counteract it. (d) We propose directions for future study on this topic.

2 Proposed Method: Voice Identity Morphing

Voice Identity Morphing (VIM), as shown in Figure 1, has two stages: a) synthetic voice generation and b) morph attack on a speaker recognition system. In the first stage, the proposed method generates synthetic speech samples exhibiting speaker-dependent speech characteristics pertaining to two different speakers, also referred to as the target speaker pair. The synthetic speech sample, called the morphed speech sample, is then compared to

³ Note that *voice morphing* as defined in this work is different from previous use of this terminology in the speech literature, where it denotes modifying an individual's voice to sound like another individual.

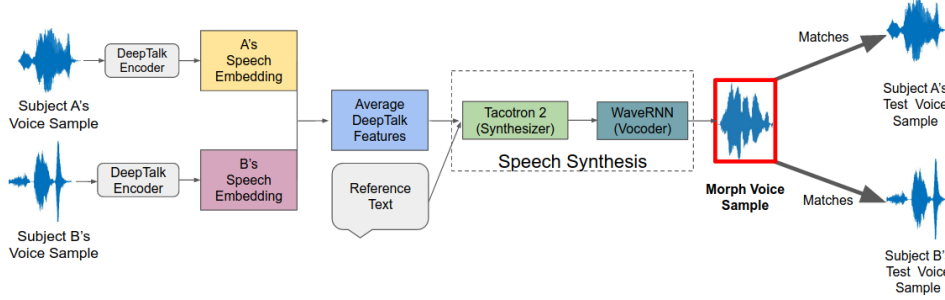


Fig. 1: Illustration of Voice Identity Morphing: Initially, the DeepTalk encoder processes and extracts embeddings that capture the unique speaker characteristics of two distinct individuals. Subsequently, to create a morphed identity, we compute the average of these two embeddings. This averaged embedding then serves as a reference point for our speech synthesis module. Ultimately, employing this reference, the vocoder module generates a spectrogram that merges elements from both contributing speakers.

individual voice samples from the target speaker pair to launch the morph attack. An attack is successful if the morphed speech sample matches both the target speaker pair’s speech samples. The morph voice generation architecture has three separate modules: Encoder, Synthesizer, and Vocoder.

We use a pre-trained **DeepTalk encoder** model to generate vocal style-based speaker identity embeddings of voice samples. We choose this encoder for its competitive performance with the x-vector system and its robustness to degraded audio scenarios. This encoder architecture consists of a 1D-CNN based speech filter bank also known as DeepVOX network [CR20] and Global Style Token (GST) [Wa18] based prosody embedding network. The DeepVOX network generates short-term speaker-dependent DeepVOX features (see Table 1 for architecture details). The GST based prosody embedding network generates a fixed dimensional reference embedding from DeepVOX features by using a 2D-CNN followed by a 128-unit GRU. The DeepTalk encoder is pre-trained on the Librispeech, VoxCeleb1 and VoxCeleb2 datasets. The synthesizer module uses these embeddings as an input during the morph sample generation stages (speech synthesis and vocoding). As an initial step of the morph sample generation stage, we average the embeddings (Emb_a and Emb_b) of two voice samples from separate speakers to generate a morph embedding $Emb_{morph} = (Emb_a + Emb_b)/2$. We perform this averaging step to incorporate features from both constituent identities. This assumes that there is an underlying geometric relationship between the identities in the learned embedding space from the DeepTalk encoder. We illustrate these relationships using t-SNE in Figure 3.

We use **Tacotron 2 speech synthesizer** [Sh18] to generate a mel-spectrogram for the corresponding text input. We use the Tacotron 2 synthesizer to retain consistency with the original DeepTalk architecture. Tacotron 2 architecture consists of an encoder and a decoder with an attention mechanism. The encoder creates an internal representation of input text, and the decoder uses the internal representation to generate features that encode the audio as a frame-level mel-spectrogram. The attention mechanism helps the decoder

Tab. 1: DeepVOX network setup for learning a 40-dimension feature representation from speech frames. All rows are convolutional layers separated by a SELU activation function.

<u>In Channels</u>	<u>Out Channels</u>	<u>Kernel</u>	<u>Dilation</u>
1	2	5x1	2x1
2	4	5x1	2x1
4	8	7x1	3x1
8	16	9x1	4x1
16	32	11x1	5x1
32	40	11x1	5x1

learn from the internal representation by weighting out potential failure cases where some subsequences of text are repeated or ignored by the decoder.

We use a **WaveRNN-based neural vocoder** [Ka18] pretrained model to generate morph samples by inverting the mel-spectrogram output from the Tacotron 2 synthesizer into audio samples. WaveRNN aims to have an expressive and non-linear transformation of the context and minimize the number of operations each step. An RNN addresses this purpose by combining the context and input within a single transformation.

3 Experimental Protocol

3.1 Dataset

We conducted experiments using the publicly accessible Librispeech dataset [Pa15], an audiobook corpus derived from Librivox projects. This dataset includes 1000 hours of audio data, in which, for each sample, a speaker reads English text. The dataset is divided into three subsets (100hr, 360hr, 500hr), all sampled at 16kHz. For our experiment, we utilized the 500 hour subset that consists of 1,166 participants (554 female and 612 male). We selected the 500hr subset not only because it is the largest subset, but also because it encompasses 440 speakers, each with more than 30 minutes of speaking time – a factor crucial for the morph generation process.

3.2 Baseline Recognition Performance

We assess speaker recognition systems’ vulnerability to morph samples using two popular speaker recognition systems: x-vector [Sn18] and ECAPA-TDNN [DTD]. We choose these systems as they are freely available and are used in a wide range of systems.⁴ We use the implementation of these systems in Speechbrain [RPO21] toolkit. The x-vector matcher is a TDNN (Time delay neural network) architecture and applies statistical pooling to extract 512-dimensional embedding for variable length utterances. The matcher utilizes categorical cross-entropy loss for training. The ECAPA-TDNN matcher architecture

⁴ ECAPA-TDNN amassed 553,704 downloads in one month (June 2023) according to the HuggingFace website [In23]

consists of convolutional layers, residual blocks, and attentive statistical pooling layers. It utilizes Additive Margin SoftMax Loss to generate a 192-dimensional embedding. Both matchers utilize Voxceleb1 [NCZ17] and Voxceleb2 [CNZ18] datasets to train the models. They use cosine distance similarity of speaker embeddings to compare a pair of speaker identities.

Before assessing their vulnerability, we evaluate the baseline recognition performance of these speaker recognition systems on 440 subjects in the 500-hr subset of the Librispeech dataset [Pa15]. Table 2 provides the performance of these speaker recognition systems in terms of True Match Rate (TMR) at 1%, 0.1%, and 0.01% False Match Rate (FMR). TMR is the proportion of genuine samples that were correctly matched, whereas FMR was the proportion of impostor samples that were incorrectly matched. ECAPA-TDNN model performs better than the x-vector model in correctly classifying genuine and impostor pairs.

Tab. 2: Performance of two speaker recognition systems in terms of TMR (%) at 1%, 0.1%, and 0.01% FMR in the Librispeech dataset. The ECAPA-TDNN and x-vector are two popular, high-performing speaker recognition systems available in the Speechbrain toolkit.

Matcher	TMR (%)		
	FMR 1%	FMR 0.1%	FMR 0.01%
ECAPA-TDNN	98.91	97.50	93.25
x-vector	88.17	78.57	68.52

3.3 Morph Generation Setup and Results

To generate morph voice samples that incorporate both identities of two different speakers, we first fine-tune a separate Tacotron 2 synthesizer with speech samples of that speaker pair. A pre-trained Tacotron 2 synthesizer needs approximately 30 minutes of the voice samples for fine-tuning [CRD21]. Therefore, we select 440 speakers (221 female and 219 male) which has 30 minutes or more cumulative duration of voice samples. From 440 speakers, we generate 96,580 speaker pairs ($^{440}C_2$). To generate *better quality* morph samples, we consider those speaker pairs which have *high similarity* in their speech. Each instance of Tacotron 2 takes 8-10 hours to fine-tune. Given this, we select the top 100 speaker pairs. We measure the similarity by the cosine distance of their ECAPA-TDNN-extracted speaker embeddings. Through this process, we select the top 100 speaker pairs, out of which only 43 pairs have unique speakers. The trimmed list of speaker pairs has 3 cross gender speaker pairs. Considering these 43 speaker pairs, we fine-tune 43 different Tacotron 2 synthesizers in parallel. For fine-tuning these Tacotron 2 synthesizers, we also provide 256-dimensional speaker embeddings extracted from a pre-trained DeepTalk encoder model [CRD21] as input along with a reference text. The fine-tuned Tacotron 2 synthesizer outputs a morphed mel spectrogram which is then fed as input into the WaveRNN vocoder [Ka18] to generate morphed speech samples. We create 100 such morphed samples from each speaker pair (10 samples per speaker) which results in 4,300 morphed samples. The speech samples used to generate morph samples are different from the ones

used for training the Tacotron 2 synthesizer. We use the remaining voice samples of a speaker for testing. Our experiment has disjoint sets of training (60%), morph (10%) and test (30%) speech samples.

To evaluate the vulnerabilities of the two speaker recognition systems against the generated morph samples (morph attack), we use the Mated Morph Presentation Match Rate (MMPMR) [Sc17] and Morphing Attack Potential (MAP) [Fe22] metrics. MMPMR is a fraction of successful morph attacks out of the total number of morph attacks. A morph attack is considered successful when the morph sample matches with test samples of both speakers. Table 3 provides the performance of morph attacks in terms of MMPMR at different thresholds corresponding to 1%, 0.1%, and 0.01% FMRs. We report the morph attack success rate in two categories: speaker pair level and morph sample level. A successful morph attack at the speaker pair level has at least one morph sample that matches the samples of both speakers. However, morph sample-level MMPMR reports the success of all morph samples irrespective of the speaker. The proposed morphing technique VIM can create morph samples attacking ECAPA-TDNN and x-vector speaker recognition systems with 95.34% and 86.04% respective success rates at 0.1% FMR, for speaker pair level. The results show that the ECAPA-TDNN speaker recognition system is more susceptible to morph attacks compared to the x-vector recognition system. The considerable success rate of morph attacks could likely be related to the morph pair selection process or the effective capturing of subject information by the DeepTalk encoding method. This infers that prior knowledge of the speaker recognition system would generate stronger morph attacks. Also, we hypothesize that state-of-the-art speaker recognition systems are likely to detect vocal features of both the parent speakers in a composite audio. This may make them vulnerable to such morphing attacks as well. We find that the fusion of speech synthesis embeddings generates effective morph audio samples for use in attacks on speaker recognition systems.

Tab. 3: Vulnerability assessment of two speaker recognition systems to voice identity morph attack in terms of MMPMR (%) at different threshold corresponding to 1%, 0.1%, and 0.01% FMR on the Librispeech dataset.

Matcher	Speaker pair MMPMR (%)			Morph sample MMPMR (%)		
	FMR 1%	FMR 0.1%	FMR 0.01%	FMR 1%	FMR 0.1%	FMR 0.01%
ECAPA-TDNN	100.00	95.34	81.39	91.23	62.11	21.58
x-vector	93.02	86.04	9.30	82.13	38.95	4.32

3.4 Result Analysis

We further analyze our morph attack performance using: 1) histogram plots, 2) t-SNE plots, and 3) morphing attack potential (MAP). Figure 2 shows the histogram plots of match scores corresponding to genuine pairs (green), impostor pairs (red), and pairs which include at least one morphed sample (blue) for both speaker recognition systems. In both systems, we find that the morphed pairs match score distribution lies between genuine and impostor score distributions. Morph samples are classified as genuine matches in the

Tab. 4: Morphing Attack Potential (MAP) [Fe22]: This metric represents the success rate (%) of a morphed sample matching at least a specified number of probe voice samples (denoted as # of attempts) within the Librispeech dataset, using one or both of the speaker recognition systems (SRS), namely ECAPA-TDNN and x-vector. The success rate is evaluated at three false match rate (FMR) thresholds: 1%, 0.1%, and 0.01%.

# of Attempts	FMR 1%		FMR 0.1%		FMR 0.01%	
	1 SRS	2 SRS	1 SRS	2 SRS	1 SRS	2 SRS
1	92.0%	52.7%	60.4%	7.6%	20.2%	2.3%
2	90.2%	46.3%	54.4%	5.7%	16.5%	1.7%
3	88.9%	41.6%	50.8%	5.0%	14.3%	1.0%
4	87.9%	38.1%	47.9%	4.6%	13.0%	0.6%
5	87.0%	35.7%	45.8%	4.2%	11.4%	0.3%

ECAPA-TDNN and x-vector systems with recognition thresholds of 0.46 and 0.96 respectively at 0.1% FMR.




Fig. 2: Speaker recognition match score distributions of non-morph versus non-morph genuine (Green), non-morph versus non-morph impostor (Red), and morph versus non-morph genuine morph scores (Blue) using ECAPA-TDNN and x-vector embeddings.

The second analysis we perform is based on the t-SNE dimensionality reduction technique. The t-SNE [VdMH08] method helps visualize high-dimensional embeddings in a two-dimensional space by reducing the dimension. Figure 3 shows the t-SNE plot of morph sample embeddings from two speaker pairs (AB and CD) along with non-morph samples of four constituent speakers (A, B, C, and D). The embeddings are extracted by the ECAPA-TDNN recognition system. Here, embeddings of morph samples of one speaker pair (AB) are closer to embeddings of A and B speakers. Similarly, embeddings of morph samples of another speaker pair (CD) are closer to embeddings of C and D speakers. The analysis again validates the effectiveness of the proposed morphing technique and the potential threat of morph attacks.




Fig. 3: t-SNE plot which illustrates high-dimensional ECAPA-TDNN embeddings of morph speech samples from two separate speaker pairs (AB and CD) and non-morph speech samples of individual speakers (A, B, C, and D). Morph embeddings of each pair are closer to the non-morph embeddings of their constituent speakers.

The Morphing Attack Potential (MAP) [Fe22] constitutes the third analysis. This metric takes into account multiple Speaker Recognition Systems (SRS) to ensure generality, and a variable number of verified probe samples for robustness. The result is a matrix (Table 4) in which one axis represents the number of probe samples (referred to as the number of attempts), and the other axis represents the number of SRS. The entries in each row represent the success rate (in percentage) of a morphed sample matching at least a specified number of probe voice samples (referred to as the number of attempts) using either or both of the SRS, viz., ECAPA-TDNN and x-vector. We report the success rates over three FMR thresholds of 1%, 0.1%, and 0.01%. The results imply that VIM is effective at a fairly competitive FMR of 1%, but suggest there is still room for improvement in performance at very low FMR thresholds. This may be attributed to the morph selection process or perhaps to the pre-trained models used in the encoder and speech synthesis steps.

4 Summary and Future Work

To the best of our knowledge, this preliminary work is the first to demonstrate the vulnerability of speaker recognition systems to morph attacks. In this regard, we propose a voice morphing technique called VIM to generate speech samples corresponding to the identities of two subjects. Using these morph samples, we demonstrate a morph attack success rate of over 80% on two popular speaker recognition systems (ECAPA-TDNN and x-vector). As future work, we propose to select high-similarity pairs for a morphing attack using x-vector to investigate whether the selection process plays a vital role in the performance of such an attack. Additionally, evaluating newer speaker recognition systems such as TitaNet [KPG22] and MFA-Conformer [Zh22] would provide more insight into the generalizability of VIM. Comparing other speech synthesis systems in the speech synthesis step would shed light on the role this step plays in the VIM attack. Furthermore, we aim to develop a

system for detecting morphed speech samples, possibly through the identification of their constituent identities. It may also be interesting to explore the maximum number of identities that can be combined into a single audio sample using VIM.

5 Reproducibility

The code for generating VIM samples can be found online at our Github link.⁵

References

- [CNZ18] Chung, J. S.; Nagrani, A.; Zisserman, A.: VoxCeleb2: Deep Speaker Recognition. In: INTERSPEECH. 2018.
- [CR20] Chowdhury, Anurag; Ross, Arun: DeepVOX: Discovering Features From Raw Audio For Speaker Recognition in Non-Ideal Audio Signals. arXiv preprint arXiv:2008.11668, 2020.
- [CRD21] Chowdhury, Anurag; Ross, Arun; David, Prabu: DeepTalk: Vocal Style Encoding for Speaker Recognition and Speech Synthesis. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 6189–6193, 2021.
- [Da18] Damer, Naser; Saladie, Alexandra Mosegui; Braun, Andreas; Kuijper, Arjan: MorGAN: Recognition Vulnerability and Attack Detectability of Face Morphing Attacks Created by Generative Adversarial Network. In: IEEE 9th international conference on biometrics theory, applications and systems (BTAS). pp. 1–10, 2018.
- [DTD] Desplanques, Brecht; Thienpondt, Jenthe; Demuynck, Kris: ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification. In: INTERSPEECH 2020. pp. 3830–3834.
- [FCM16] Ferrara, Matteo; Cappelli, Raffaele; Maltoni, Davide: On The Feasibility of Creating Double-Identity Fingerprints. IEEE Transactions on Information Forensics and Security, 12(4):892–900, 2016.
- [Fe22] Ferrara, Matteo; Franco, Annalisa; Maltoni, Davide; Busch, Christoph: Morphing Attack Potential. In: International Workshop on Biometrics and Forensics (IWBF). IEEE, pp. 1–6, 2022.
- [FFM14] Ferrara, Matteo; Franco, Annalisa; Maltoni, Davide: The Magic Passport. In: IEEE International Joint Conference on Biometrics. pp. 1–7, 2014.
- [Ho18] Hoy, Matthew B: Alexa, Siri, Cortana, and More: An Introduction to Voice Assistants. Medical Reference Services Quarterly, 37(1):81–88, 2018.
- [In23] Inc., Hugging Face: , Hugging Face: The AI community building the future, 2023. Accessed: 2023-07-08.
- [JFR07] Jain, Anil K; Flynn, Patrick; Ross, Arun A: Handbook of Biometrics. Springer Science & Business Media, 2007.

⁵ <https://github.com/morganlee123/VIM>

- [Ka18] Kalchbrenner, Nal; Elsen, Erich; Simonyan, Karen; Others: Efficient Neural Audio Synthesis. In: International Conference on Machine Learning. PMLR, pp. 2410–2419, 2018.
- [KPG22] Koluguri, Nithin Rao; Park, Taejin; Ginsburg, Boris: TitaNet: Neural Model for Speaker Representation with 1D Depth-Wise Separable Convolutions and Global Context. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 8102–8106, 2022.
- [MFFM19] Matteo Ferrara, Matteo; Franco, Annalisa; Maltoni, Davide: Decoupling Texture Blending and Shape Warping in Face Morphing. In: Proceedings of the 18th International Conference of the Biometrics Special Interest Group (BIOSIG). Gesellschaft für Informatik eV, 2019.
- [MSI01] Melin, Håkan; Sandell, Anna; Ihse, Magnus: CTT-Bank: A Speech Controlled Telephone Banking System-An Initial Evaluation. TMH-QPSR, 1:1–27, 2001.
- [NCZ17] Nagrani, A.; Chung, J. S.; Zisserman, A.: VoxCeleb: A Large-Scale Speaker Identification Dataset. In: INTERSPEECH. 2017.
- [Ng22] Nguyen, Huy H.; Marcel, Sebastien; Yamagishi, Junichi; Echizen, Isao: Master Face Attacks on Face Recognition Systems. IEEE Transactions on Biometrics, Behavior, and Identity Science, 4(3):398–411, 2022.
- [Pa15] Panayotov, Vassil; Chen, Guoguo; Povey, Daniel; Khudanpur, Sanjeev: Librispeech: An ASR Corpus Based on Public Domain Audio Books. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 5206–5210, 2015.
- [RB17] Rathgeb, Christian; Busch, Christoph: On The Feasibility of Creating Morphed Iris Codes. In: IEEE International Joint Conference on Biometrics (IJCB). pp. 152–157, 2017.
- [RMR17] Roy, Aditi; Memon, Nasir; Ross, Arun: MasterPrint: Exploring the Vulnerability of Partial Fingerprint-Based Authentication Systems. IEEE Transactions on Information Forensics and Security, 12(9):2013–2025, 2017.
- [RPO21] Ravanelli, Mirco; Parcollet, Titouan; Others: , SpeechBrain: A General-Purpose Speech Toolkit, 2021. arXiv:2106.04624.
- [RRB16] Ramachandra, R; Raja, KB; Busch, C: Detecting Morphed Face Images. In: 8th IEEE International Conference on Biometrics Theory, Applications and Systems, BTAS. pp. 1–7, 2016.
- [Sc17] Scherhag, Ulrich; Nautsch, Andreas; Rathgeb, Christian; Others: Biometric Systems Under Morphing Attacks: Assessment of Morphing Techniques and Vulnerability Reporting. In: International Conference of the Biometrics Special Interest Group (BIOSIG). IEEE, pp. 1–7, 2017.
- [Sh18] Shen, Jonathan; Pang, Ruoming; Weiss, Ron J; Others: Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 4779–4783, 2018.
- [Sn18] Snyder, David; Garcia-Romero, Daniel; Sell, Gregory; Povey, Daniel; Khudanpur, Sanjeev: X-Vectors: Robust DNN Embeddings for Speaker Recognition. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 5329–5333, 2018.

- [SR21] Sharma, Renu; Ross, Arun: Image-Level Iris Morph Attack. In: IEEE International Conference on Image Processing (ICIP). pp. 3013–3017, 2021.
- [VdMH08] Van der Maaten, Laurens; Hinton, Geoffrey: Visualizing Data Using t-SNE. *Journal of Machine Learning Research*, 9(11), 2008.
- [Ve21] Venkatesh, Sushma; Ramachandra, Raghavendra; Raja, Kiran; Busch, Christoph: Face Morphing Attack Generation and Detection: A Comprehensive Survey. *IEEE Transactions on Technology and Society*, 2(3):128–145, 2021.
- [Wa18] Wang, Yuxuan; Stanton, Daisy; Zhang, Yu; Others: Style Tokens: Unsupervised Style Modeling, Control and Transfer in End-to-End Speech Synthesis. In: International Conference on Machine Learning (ICML). PMLR, pp. 5180–5189, 2018.
- [Zh21] Zhang, Haoyu; Venkatesh, Sushma; Ramachandra, Raghavendra; Others: MIP-GAN—Generating Strong and High Quality Morphing Attacks Using Identity Prior Driven GAN. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(3):365–383, 2021.
- [Zh22] Zhang, Yang; Lv, Zhiqiang; Wu, Haibin; Others: MFA-Conformer: Multi-scale Feature Aggregation Conformer for Automatic Speaker Verification. In: INTERSPEECH. 2022.