

SegReConcat: A Data Augmentation Method for Voice Anonymization Attack

Ridwan Arefeen*, Xiaoxiao Miao[†], Rong Tong*, Aik Beng Ng[‡], Simon See[‡]

* Singapore Institute of Technology

E-mail: 2403754@sit.singaporetech.edu.sg, tong.rong@singaporetech.edu.sg

[†] Duke Kunshan University, China

E-mail: xiaoxiao.miao@dukekunshan.edu.cn

[‡] NVIDIA AI Technology Centre

E-mail: {aikbengn,ssee}@nvidia.com

Abstract—Anonymization of voice seeks to conceal the identity of the speaker while maintaining the utility of speech data. However, residual speaker cues often persist, which pose privacy risks. We propose SegReConcat, a data augmentation method for attacker-side enhancement of automatic speaker verification systems. SegReConcat segments anonymized speech at the word level, rearranges segments using random or similarity-based strategies to disrupt long-term contextual cues, and concatenates them with the original utterance, allowing an attacker to learn source speaker traits from multiple perspectives. The proposed method has been evaluated in the VoicePrivacy Attacker Challenge 2024 framework across seven anonymization systems, SegReConcat improves de-anonymization on five out of seven systems¹.

I. INTRODUCTION

Voice data inherently contains rich personal information, including speaker identity, emotional state, and potential medical conditions. When raw speech is shared, for example, with cloud services or on social media, adversaries can exploit these biometric cues to conduct identity theft or enable targeted advertising [1]. As a result, voice privacy has become a critical research objective, often viewed as an ongoing game between users (defenders) and attackers (adversaries).

On one hand, users employ voice anonymization or de-identification techniques to modify original speech utterances to conceal the speaker's identity while preserving linguistic content and permissible paralinguistic attributes. This process generates anonymized speech that can be safely shared. The VoicePrivacy Challenges (VPC) have played a key role in advancing this field by providing benchmarks [2], [3], [4], [5] and fostering the development of anonymization technologies [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22].

On the other hand, attackers leverage their knowledge of anonymization systems to infer residual identifying information from anonymized speech. While most research focuses on improving anonymization from the users' perspective and assumes fixed attack scenarios, where attackers primarily rely on

automatic speaker verification (ASV) systems such as ECAPA-TDNN [23] to trace source speaker information, relatively little attention has been paid to developing comprehensive methods for evaluating the robustness of anonymization against determined adversaries. To facilitate research in this area, the same committee that organized the VPCs launched the first VoicePrivacy Attacker Challenge (VPAC) in 2024 [24], [25]. The challenge provided anonymized datasets generated by various anonymization systems [5], along with a baseline attacker system. Participants developed ASV systems to evaluate the robustness of different anonymization methods. The top submissions primarily improved upon three aspects: (i) Data augmentation: the SpecWav attack [26], combining wav2vec2.0 features [27] with spectrogram resizing data augmentation technique. (ii) Neural network-based embedding extraction: LoRA-adapted ResNet34 [28], [29], [30], as well as fine-tuned the TitaNet-Large model [31]. (iii) Backend models: [32] utilize ECAPA-TDNN feature extractor trained on mixed datasets with a PLDA-based [33], [34] scoring backend trained on anonymized data, incorporating SpecAugment [35] for additional robustness. In addition, other submissions explored alternative approaches, such as using different distance metrics and speaker embedding normalization techniques [36].

Although the attack strategies described above have proven effective in exposing weaknesses in anonymization systems [25], certain domains remain underexplored. Recent studies have shown that identity-revealing cues may persist in prosodic, phonetic, and linguistic content features that remain unchanged after anonymization [37], [20], [15], [16], also in speech temporal dynamics [38], [39]. For instance, both linguistic content and prosody can inadvertently preserve speaker-specific characteristics in the anonymized output [7]. Moreover, inference attacks that match pseudo-speaker embeddings against a known speaker pool have been shown to outperform generic ASV models in de-anonymization tasks [40].

Motivated by these findings, this paper follows the VPAC 2024 framework and proposes a novel attacker-side data augmentation technique, called **SegReConcat**, where an anonymized speech utterance is **segmented** into word segments, then **rearranges** the words according to different strate-

¹Code can be found at https://github.com/monkeyDarefeen/SegReConcat_augmenter. Xiaoxiao Miao is the corresponding author and this work was conducted while she was at SIT.

gies, and finally **concatenates** with the input anonymized speech. The rearranged anonymized speech is generated by permuting the word sequence of the anonymized utterance².

Since both content and prosodic information can reveal speaker identity potentially in both short-term and long-term contexts, an ASV model trained on anonymized speech with the original word order may inadvertently learn residual speaker traits from long-term sequential dependencies, which are easier to capture. By disrupting the word order, **SegReConcat** breaks these long-term contextual cues, thereby forcing the model to focus on extracting residual speaker information from short-term, word-level features. This makes the attack more targeted and potentially more effective in revealing subtle identity cues. Additionally, by concatenating the rearranged speech with the original anonymized speech to create augmented inputs, the model is enabled to learn source speaker traits from multiple perspectives.

We explore different strategies for word rearrangement, including randomly shuffling words and grouping similar words based on similarity scores. We evaluated our attack method and demonstrated consistent reductions in equal error rate (EER) across 5 out of 7 anonymization systems provided by VPAC 2024.

II. RELATED WORK

In this section, we present the official design of the VPAC 2024, which provides the framework for this study. We also review existing attacker approaches, with a focus on data augmentation techniques in this domain and their limitations.

A. The VoicePrivacy Attacker Challenge

The VPAC 2024 [24], [25] offers a standardized framework to evaluate voice anonymization systems under a semi-informed attacker model. Attackers build ASV systems using anonymized speech to determine whether two utterances belong to the same speaker, aiming to lower EER, lower EER indicates a more effective recovery of speaker identity.

The attacker evaluates anonymized LibriSpeech [41] corpus for seven anonymization systems, each system applies a different anonymization strategy, such as disentanglement-based speaker embedding manipulation and codec-based synthesis. These systems include three VPC 2024 baselines (*B3–B5*) and four top submissions (*T8-5*, *T10-2*, *T12-5*, *T25-1*) [5].

- *B3* [42]: phonetic transcription with pitch/energy modification and generated pseudo-speaker embeddings
- *B4* [7]: a neural audio codec with language modeling of semantic tokens
- *B5* [6]: vector-quantized bottleneck (VQ-BN) features from an ASR model with original pitch.

²Note that the anonymized speech provided by VPAC 2024 always retains the original word sequence to preserve utility, such as content and emotion, which are important to users, while still ensuring strong concealment of the original speaker’s identity. However, as an attacker, one can reorder the sequence to exploit residual speaker identity information.

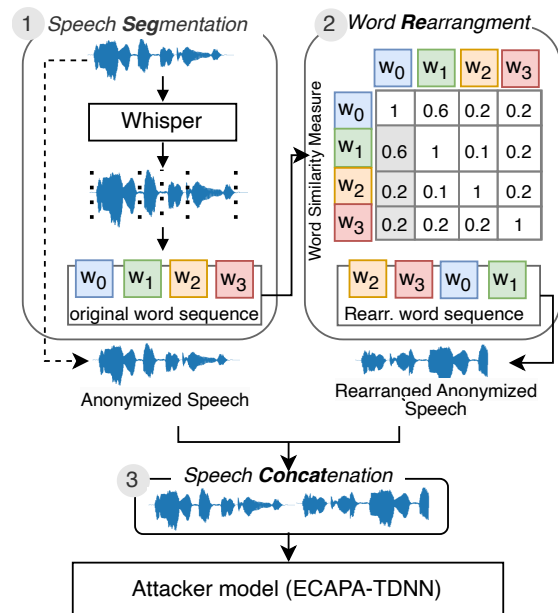


Fig. 1. (1) SegReConcat Flow Diagram

- *T8-5* [15]: a hybrid scheme randomly choosing between (1) ASR(Whisper)+TTS(VITS) and (2) k-NN voice conversion on WavLM features.
- *T10-2* [43]: a neural codec (using VQ) that explicitly disentangles linguistic content, speaker identity, and emotion.
- *T12-5* [44]: an extension of B5 with additional pitch smoothing.
- *T25-1* [45]: VQ-BN content features with global style tokens and emotion transfer for speaker style.

Each system’s anonymized data includes both training and evaluation subsets. The training set LibriSpeech *train-clean-360* contains approximately 104,014 anonymized utterances derived from 921 speakers. The development and test sets are anonymized versions of LibriSpeech *dev-clean* and *test-clean* respectively, each comprising both female-female and male-male speaker trial pairs.

B. Existing Attacker Approaches

Attacker systems have been developed from various perspectives, as outlined in Section I. This section primarily reviews attacker-side data augmentation techniques, particularly those involving modifications to acoustic features that can be applied to many attacker systems. These include methods such as SpecAugment [46], which randomly masks time frames or frequency bands. One system submitted to the VPAC 2024 employed SpecAugment alongside a mixed-data training strategy, combining anonymized and original speech to improve speaker discriminability. Additionally, techniques that resize the spectrogram by stretching or compressing it along the frequency axis have been proposed for attacker models [26]. After stretching, the excess high-frequency components are

trimmed; after compression, zero-padding is applied to the high-frequency band to maintain consistent frequency dimensions before and after resizing.

The motivation behind these methods is to introduce perturbations into the spectrogram, forcing the attacker model to learn speaker-relevant features that are robust to such distortions. Unlike existing data augmentation techniques, this work proposes **SegReConcat**, a perceptually motivated data augmentation method. Details follow in the next section.

III. PROPOSED SEGRECONCAT DATA AUGMENTATION

This section details the proposed **SegReConcat** framework, which consists of three stages: **Segmentation**, **Rearrangement**, and **Concatenation**, as shown in Figure 1. It takes an anonymized utterance, segments it at the word level, rearranges the words, and finally concatenates the rearranged sequence with the original audio.

The motivation for segmenting and rearranging the anonymized word sequence is to break the continuous content flow of speech and thereby disrupt the long-term temporal dynamics of speaker information. We assume that certain prosodic and phonetic patterns, such as coarticulation style, pitch dynamics, and spectral characteristics, remain speaker-dependent even after anonymization. By reordering word-level segments, we aim to amplify these residual stylistic cues in a way that is consistent for each speaker.

The motivation for combining rearranged speech with its original anonymized version within a single training instance, so the ASV model observes different realizations of the same linguistic content. This encourages learning speaker traits invariant across different segment orders, promoting reliance on speaker-specific acoustic patterns rather than content structure. Although data volume doubles, the structural correlation introduced enhances speaker-discriminative representation learning.

We now elaborate on each step of the proposed **SegReConcat** method.

A. Segmentation

Given an anonymized utterance u , the first step is to perform word-level segmentation using the Whisper-medium ASR model³, which achieves a 2.9% word error rate (WER) on LibriSpeech *test-clean* [47]. This high accuracy on the LibriSpeech dataset ensures reliable word boundary detection and decompose the utterance into N discrete word segments represented as $\mathbf{w}^o = w_1, w_2, \dots, w_N$.

B. Rearrangement

Next, we apply a rearrangement function g to the sequence \mathbf{w}^o , a permuted word order $\mathbf{w}^r = g(\mathbf{w}^o)$. The aim is to disrupt the original content flow and long-range temporal dependencies that may preserve speaker characteristics, even after anonymization. The function g processes each utterance individually. Three rearrangement strategies are introduced. One straightforward method is to randomly shuffle the original

sequence \mathbf{w}^o to remove natural sequential cues. This is referred to as random rearrangement (**RR**).

The other two strategies are similarity-based approaches, aiming to maximize the total similarity between consecutive words in \mathbf{w}^r . This aims to investigate whether grouping similar words can amplify speaker style characteristics, thereby enhancing source speaker learning in ASV model. To measure word similarity, we consider both acoustic and semantic features.

1) *Acoustic Feature-Based Rearrangement (AR)*: Let us denote the MFCC feature matrices for two word segments w_i and w_j as $\mathbf{M}_i \in \mathbb{R}^{d \times T_i}$ and $\mathbf{M}_j \in \mathbb{R}^{d \times T_j}$, respectively, where $d = 13$ is the MFCC dimension, and T_i and T_j are the frame lengths of the corresponding word segments. Considering the different frame lengths across word segments, the Dynamic Time Warping (DTW) distance [48], using Euclidean frame-wise cost, is computed to measure the acoustic similarity between segments, and is denoted as $D_{ij} = \text{DTW}(\mathbf{M}_i, \mathbf{M}_j)$, and is then converted into a similarity score via $s_{ij} = \frac{1}{1+D_{ij}}$.

To obtain a \mathbf{w}^r that optimizes the total similarity between consecutive words, we perform a greedy traversal: starting from the first segment (or an initial silence segment, if detected), we repeatedly append the unused segment with the highest similarity to the current one. Any remaining segments (if not reachable) are appended in arbitrary order.

2) *Semantic Feature-Based Rearrangement (SR)*: Besides using MFCC features to measure word similarity, we explore semantic feature measurements, as they are expected to carry more content information and better group words with similar pronunciations. Each segment is passed through the encoder of the same pretrained Whisper-medium ASR model used for utterance segmentation in Section III-A to obtain a hidden segment-level representation $\mathbf{h} \in \mathbb{R}^{2 \times d'}$, obtained by concatenating the mean and max pooled frame-level features over time, where d' is the dimension of the last layer of the encoder, i.e. 1024.

Pairwise word segment w_i and w_j similarities are computed using cosine similarity: $h_{ij} = \frac{\mathbf{h}_i^T \mathbf{h}_j}{\|\mathbf{h}_i\| \|\mathbf{h}_j\|}$. Similar to **AR**, \mathbf{w}^r is constructed by selecting the segment with the highest average similarity as the starting point, and then iteratively appending the most similar remaining segment at each step.

C. Concatenation

To further enhance the diversity and robustness, we concatenate the original \mathbf{w}^o and the rearranged sequence \mathbf{w}^r , forming an augmented sequence $\mathbf{w}^{\text{aug}} = \text{concat}(\mathbf{w}^o, \mathbf{w}^r)$, totaling $2N$ segments.

IV. EXPERIMENTS

To evaluate the effectiveness of the proposed **SegReConcat** method, we conducted systematic experiments within the framework of the VPAC 2024. We begin by focusing on one anonymization system, *T8-5*, selected for its relatively strong privacy performance and architectural diversity. This serves as our development benchmark for investigating various configurations of the proposed method, including different segment

³<https://huggingface.co/openai/whisper-medium>

TABLE I
THE TABLE REPORTS AVERAGED EER ON T8-5 ANONYMIZATION SYSTEM.

Methods	Avg EER (%)
No Augmentation	37.39
SpecAug	37.93
RR	38.36
+ Concatenation	26.51
+ Concatenation + SpecAug	25.94
AR	38.34
+ Concatenation	26.80
+ Concatenation + SpecAug	27.82
SR	39.10
+ Concatenation	27.42
+ Concatenation + SpecAug	32.70

reordering strategies (e.g., random shuffle **RR**, acoustic-based **AR**, and semantic-based **SR**, with and without concatenation, and in combination with additional spectral augmentation.

A. System Configurations

We fully follow the VPAC 2024 [24] protocol as previously described in Section II-A. and experiment with the proposed data augmentation method using different configurations on all seven anonymization systems. All experiments share the same speaker verification back-end, which is based on the ECAPA-TDNN architecture with 1,024 channels in the convolution frame layers and implemented via SpeechBrain⁴. Speaker embeddings are extracted from the final layer, and similarity scoring is performed using cosine distance between averaged enrollment and test embeddings. The model is trained from scratch using anonymized training data only. Regarding the training hyperparameters, we primarily follow the default settings in SpeechBrain, however, unlike the default of 10 epochs, we train the model until the development loss stops decreasing and select the model that achieves the best development loss. EER is used as the primary metric for assessing re-identification performance. A lower EER indicates a stronger attack.

B. Experimental Results and Analysis

1) *Ablation Study on T8-5*: Table I compares the EER results on T8-5 without data augmentation, with the commonly used SpecAug, along with our proposed **SegReConcat** under different configurations. The first observation is that SpecAug achieves a very similar EER to the baseline system without any data augmentation. And using rearrangement data alone, regardless of which strategy, results in slightly higher EERs, indicating worse attack performance. One possible reason is that training solely on rearranged data confuses the ASV attacker system, preventing it from learning speaker-discriminative cues from the disordered sequences. However, applying the concatenation method on top of the RR, AR, or

SR strategies reduces the EER by around 10% (from 37.39% to around 27%), demonstrating an effective attacking ability.

If we compare the results of **SegReConcat** using different rearrangement approaches, the trend is that ‘RR + concatenation’ (26.51%) performs slightly better than ‘AR + concatenation’ (26.80%), with ‘SR + concatenation’ (27.42%) being slightly worse, though overall the differences are minor. Notably, the similarity-based rearrangement approaches (AR/SR) require additional computation for similarity measurement yet fail to surpass the performance of simple random arrangement (RR). This suggests that RR’s inherent sequence diversity may help prevent the ASV model from overfitting to artificially constructed acoustic patterns.

Next, considering that SpecAug masks the acoustic features while **SegReConcat** permutes the words, the two data augmentation methods are fundamentally different. It is interesting to investigate whether these augmentations are complementary. Therefore, we combine both methods (highlighted in grey) and find that the improvements are not consistent, ‘RR + Concatenation + SpecAug’ achieves the lowest EER (25.94%).

In summary, through the ablation study on T8-5, we have verified that the proposed **SegReConcat** is effective, with either the ‘RR + concatenation’ or the ‘RR + concatenation + SpecAug’ configuration performing the best. We now proceed to experiments on the other systems using these configurations.

2) *Results Across Seven Anonymized Systems*: Experimental results on all seven anonymization systems are shown in Table II, where **SegReConcat** specifically denotes the use of RR as the similarity metric during the rearrangement stage. We evaluated both **SegReConcat** alone and in combination with SpecAug for all anonymization systems, and report the best results achieved. Among the seven systems, five show lower EERs and thus stronger attacks. Of these, two achieve the lowest EER using **SegReConcat** alone, while the other three achieve the lowest EER using **SegReConcat** in combination with SpecAug. The absolute average EER reductions range from 1% to 3%, with the exception of T8-5, which achieves an average absolute EER reduction of around 11%.

For the two systems that did not show a decrease in EER after applying the proposed data augmentation, i.e., B5 and T12-5 (see system descriptions in Section II-A), it is noteworthy that both use VQ layers to quantize the continuous self-supervised learning features. A possible explanation is that our proposed data augmentation method is designed to disrupt the continuous content flow of speech and interfere with the long-term temporal dynamics of speaker information. However, the VQ process already discretizes the SSL features and removes continuity, potentially rendering **SegReConcat** less effective. Note that T25-1 also uses VQ-BN features, but **SegReConcat** still achieves a lower EER. This may be due to its additional use of emotion transfer technology, which could leak temporal speaker dynamics. These dynamics can be further disrupted by **SegReConcat**, enhancing the exploitation of speaker-specific features for the attacker’s ASV model.

⁴<https://speechbrain.readthedocs.io/>

TABLE II
PERFORMANCE COMPARISON OF THE ECAPA-TDNN ATTACKER SYSTEMS WITH AND WITHOUT SEGReConcat (USING **RR**), IN TERMS OF EER (%), AGAINST SEVEN ANONYMIZATION SYSTEMS.

Anonymization System	Attacker System	SpecAugment	EER on dev-clean subset			EER on test-clean subset			Avg. EER
			Female	Male	Avg.	Female	Male	Avg.	
B3	Baseline	✗	27.56	22.98	25.27	29.01	27.84	28.43	26.85
	SegReConcat+SpecAug	✓	27.58	22.98	25.28	28.79	24.28	26.53	25.91
B4	Baseline	✗	35.09	28.73	31.91	26.10	29.62	27.86	29.89
	SegReConcat	✗	30.67	24.99	27.83	23.54	25.17	24.36	26.09
B5	Baseline	✗	32.53	29.21	30.87	30.66	31.21	30.94	30.90
	SegReConcat+SpecAug	✓	35.62	30.24	32.93	33.25	32.52	32.89	32.91
T8-5	Baseline	✗	33.49	38.81	36.15	39.58	37.67	38.63	37.39
	SegReConcat+SpecAug	✓	26.42	26.86	26.64	25.55	24.94	25.25	25.94
T10-2	Baseline	✗	45.74	40.84	43.29	39.60	41.43	40.52	41.90
	SegReConcat+SpecAug	✓	41.62	35.25	38.44	37.59	39.20	38.40	38.41
T12-5	Baseline	✗	34.94	31.70	33.32	30.65	35.02	32.84	33.08
	SegReConcat+SpecAug	✓	35.94	33.54	34.74	34.31	36.97	35.64	35.19
T25-1	Baseline	✗	40.77	37.60	39.19	39.60	38.53	39.06	39.12
	SegReConcat	✗	38.35	35.10	36.73	37.77	35.41	36.59	36.66

V. CONCLUSIONS

This work investigates the vulnerability of the current voice anonymization systems from the attacker’s perspective by introducing a novel augmentation strategy, **SegReConcat**, which can be applied to any back-end ASV model to enhance re-identification. This study is motivated by the observation that anonymized speech often retains residual speaker-specific cues in prosody and articulation. The proposed method leverages segment-level reorganization and concatenation to amplify such cues and improve speaker discriminability. Experimental results on the *T8-5* anonymization system demonstrate that the proposed approach yields improvement in speaker verification accuracy, achieving an 11% absolute reduction in average EER compared to the baseline. Further evaluation across all seven anonymization systems in the VoicePrivacy 2024 Attacker Challenge[24] shows superior attacking performance for five systems, confirming the robustness and generalizability of the method.

These findings highlight limitations in current anonymization pipelines, particularly in their inability to fully suppress subtle speaker identity traces. Our results emphasize the need for future anonymization systems to consider attacker-informed augmentations and prosodic invariance more explicitly in their design.

VI. ACKNOWLEDGMENT

This work is supported by Ministry of Education, Singapore, under its Academic Research Tier 1 (R-R13-A405-0005). Thanks Timothy Liu from Nvidia AI Technology Centre for his insightful suggestions.

REFERENCES

[1] “General data protection regulation (GDPR),” <https://gdpr.eu/what-is-gdpr>.

[2] N. Tomashenko, X. Wang, E. Vincent, J. Patino, B. M. L. Srivastava, P.-G. Noé, A. Nautsch, N. Evans, J. Yamagishi, B. O’Brien, A. Chancelu, J.-F. Bonastre, M. Todisco, and M. Maouche, “The voiceprivacy 2020 challenge: Results and findings,” *Computer Speech & Language*, vol. 74, p. 101362, Jul. 2022.

[3] N. Tomashenko, X. Miao, P. Champion, S. Meyer, X. Wang, E. Vincent, M. Panariello, N. Evans, J. Yamagishi, and M. Todisco, “The voiceprivacy 2024 challenge evaluation plan,” *arXiv preprint arXiv:2404.02677*, 2024.

[4] M. Panariello, N. Tomashenko, X. Wang, X. Miao, P. Champion, H. Nourtel, M. Todisco, N. Evans, E. Vincent, and J. Yamagishi, “The VoicePrivacy 2022 challenge: Progress and perspectives in voice anonymisation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, pp. 1–14, 2024.

[5] N. Tomashenko, X. Miao, P. Champion, S. Meyer, X. Wang, E. Vincent, M. Panariello, N. Evans, J. Yamagishi, and M. Todisco, “The VoicePrivacy 2024 challenge evaluation plan,” *arXiv preprint arXiv:2404.02677*, 2024.

[6] P. Champion, S. Ouni, D. Jouviet, and A. Larcher, “Anonymizing speech: Evaluating and designing speaker anonymization techniques,” *arXiv preprint*, vol. arXiv:2308.04455v4, 2023.

[7] M. Panariello, F. Nespola, M. Todisco, and N. Evans, “Speaker anonymization using neural audio codec language models,” 2024.

[8] S. Meyer, F. Lux, J. Koch, P. Denisov, P. Tilli, and N. T. Vu, “Prosody is not identity: A speaker anonymization approach using prosody cloning,” in *Proc. IEEE ICASSP*. IEEE, 2023, pp. 1–5.

[9] A. S. Shamsabadi, B. M. L. Srivastava, A. Bellet, N. Vauquier, E. Vincent, M. Maouche, M. Tommasi, and N. Papernot, “Differentially private speaker anonymization,” *Proceedings on Privacy Enhancing Technologies*, vol. 2023, no. 1, Jan. 2023.

[10] C. O. Mawalim, K. Galajit, J. Karnjana, S. Kidani, and M. Unoki, “Speaker anonymization by modifying fundamental frequency and x-vector singular value,” *Computer Speech & Language*, vol. 73, p. 101326, 2022.

[11] J. Yao, Q. Wang, P. Guo, Z. Ning, Y. Yang, Y. Pan, and L. Xie, “MUSA: Multi-lingual speaker anonymization via serial disentanglement,” *arXiv preprint arXiv:2407.11629*, 2024.

[12] X. Miao, X. Wang, E. Cooper, J. Yamagishi, and N. Tomashenko, “Language-independent speaker anonymization approach using self-supervised pre-trained models,” in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2022)*, 2022, pp. 279–286.

[13] —, “Speaker anonymization using orthogonal householder neural network,” *IEEE/ACM Trans. Audio, Speech, and Language Processing*, vol. 31, pp. 3681–3695, 2023.

[14] J. Yao, Q. Wang, P. Guo, Z. Ning, and L. Xie, “Distinctive and natural speaker anonymization via singular value transformation-assisted

- matrix,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [15] H. L. Xinyuan, Z. Cai, A. Garg, K. Duh, L. P. García-Perera, S. Khudanpur, N. Andrews, and M. Wiesner, “Hltcoe jhu submission to the voice privacy challenge 2024,” *arXiv preprint arXiv:2409.08913*, 2024.
 - [16] Z. Cai, H. L. Xinyuan, A. Garg, L. P. García-Perera, K. Duh, S. Khudanpur, N. Andrews, and M. Wiesner, “Privacy versus emotion preservation trade-offs in emotion-preserving speaker anonymization,” in *2024 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2024, pp. 409–414.
 - [17] S. Ghosh, M. Jouaiti, A. Das, Y. Sinha, T. Polzehl, I. Siegert, and S. Stober, “Anonymising elderly and pathological speech: Voice conversion using ddsp and query-by-example,” *Interspeech*, 2024.
 - [18] Z. Shaohu, L. Zhouyu, and D. Anupam, “VoicePM: A robust privacy measurement on voice anonymity,” in *Proc. 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2023, p. 215–226.
 - [19] X. Miao, R. Tao, C. Zeng, and X. Wang, “A benchmark for multi-speaker anonymization,” *IEEE Transactions on Information Forensics and Security*, vol. 20, pp. 3819–3833, 2025.
 - [20] X. Miao, Y. Zhang, X. Wang, N. Tomashenko, D. C. L. Soh, and I. McLoughlin, “Adapting general disentanglement-based speaker anonymization for enhanced emotion preservation,” *Computer Speech & Language*, p. 101810, 2025.
 - [21] J. Yao, H. Liu, E. S. Chng, and L. Xie, “Easy: Emotion-aware speaker anonymization via factorized distillation,” *Accepted by Interspeech 2025*, 2025.
 - [22] C. Franzreb, A. Das, T. Polzehl, and S. Möller, “Private knn-vc: Interpretable anonymization of converted speech,” in *Accepted by Interspeech 2025*, 2025.
 - [23] B. Desplanques, J. Thienpondt, and K. Demuyne, “Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” in *Interspeech 2020*. ISCA, 2020.
 - [24] N. Tomashenko, X. Miao, E. Vincent, and J. Yamagishi, “The first voiceprivacy attacker challenge evaluation plan,” 2024.
 - [25] —, “The first voiceprivacy attacker challenge,” in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, Apr. 2025, p. 1–2.
 - [26] Y. Li, Y. Zheng, Z. Guo, Y. Wang, J. Yin, and H. Fei, “Specwav-attack: Leveraging spectrogram resizing and wav2vec 2.0 for attacking anonymized speech,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–2.
 - [27] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NIPS*, vol. 33, 2020, pp. 12 449–12 460.
 - [28] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen *et al.*, “Lora: Low-rank adaptation of large language models,” *ICLR*, vol. 1, no. 2, p. 3, 2022.
 - [29] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
 - [30] X. Lyu, Y. Wang, T. Zhao, and H. Liu, “Fast adaptation of pretrained speaker verification system for source speaker tracking,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–2.
 - [31] C. O. Mawalim, A. Adila, and M. Unoki, “Fine-tuning titanet-large model for speaker anonymization attacker systems,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–2.
 - [32] Y. Zhang, Z. Bi, F. Xiao, X. Yang, Q. Zhu, and J. Guan, “Attacking voice anonymization systems with augmented feature and speaker identity difference,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–2.
 - [33] S. J. Prince and J. H. Elder, “Probabilistic linear discriminant analysis for inferences about identity,” in *2007 IEEE 11th international conference on computer vision*. IEEE, 2007, pp. 1–8.
 - [34] S. Ioffe, “Probabilistic linear discriminant analysis,” in *European Conference on Computer Vision*. Springer, 2006, pp. 531–542.
 - [35] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Interspeech 2019*. ISCA, Sep. 2019.
 - [36] H. L. Xinyuan, A. Garg, Z. Cai, K. Duh, L. P. García-Perera, S. Khudanpur, N. Andrews, and M. Wiesner, “Hltcoe submission to the voiceprivacy attacker challenge,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–2.
 - [37] U. E. Gaznepoglu and N. Peters, “Why disentanglement-based speaker anonymization systems fail at preserving emotions?” in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
 - [38] N. Tomashenko, E. Vincent, and M. Tommasi, “Analysis of speech temporal dynamics in the context of speaker verification and voice anonymization,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
 - [39] —, “Exploiting context-dependent duration features for voice anonymization attack systems,” in *Interspeech 2025*, 2025.
 - [40] W. Bao, M. Jadhav, and V. Bindschaedler, “Inference Attacks for X-Vector Speaker Anonymization,” in *2025 IEEE Security and Privacy Workshops (SPW)*, 2025, pp. 152–159.
 - [41] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “LibriSpeech: an ASR corpus based on public domain audio books,” in *Proc. ICASSP*. IEEE, 2015, pp. 5206–5210.
 - [42] S. Meyer, P. Tilli, P. Denisov, F. Lux, J. Koch, and N. T. Vu, “Anonymizing speech with generative adversarial networks to preserve speaker privacy,” in *2022 IEEE Spoken Language Technology Workshop (SLT)*, 2023, pp. 912–919.
 - [43] J. Yao, N. Kuzmin, Q. Wang, P. Guo, Z. Ning, D. Guo, K. A. Lee, E.-S. Chng, and L. Xie, “NPU-NTU System for Voice Privacy 2024 Challenge,” *arXiv preprint arXiv:2409.04173*, 2024.
 - [44] N. Kuzmin, H.-T. Luong, J. Yao, L. Xie, and K. A. Lee, “NTU-NPU System for Voice Privacy 2024 Challenge,” *SPSC 2024*, 2024. [Online]. Available: https://www.voiceprivacychallenge.org/vp2024/docs/T12__NTU-NPU_System_for_Voice_Privacy_2024_Challenge.pdf
 - [45] W. Gu, Z. Liu, L. Chen, R. Wang, C. Guo, W. Guo, K. A. Lee, and Z.-H. Ling, “USTC-PolyU system for the VoicePrivacy 2024 Challenge,” *SPSC 2024*, 2024. [Online]. Available: https://www.voiceprivacychallenge.org/vp2024/docs/T25__USTC-PolyU_system_for_the_VoicePrivacy_2024_Challenge.pdf
 - [46] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *Interspeech 2019*, Sep 2019.
 - [47] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” 2022.
 - [48] M. Müller, “Dynamic time warping,” *Information retrieval for music and motion*, pp. 69–84, 2007.