

Comparing Speech Anonymization Efficacy by Voice Conversion Using KNN and Disentangled Speaker Feature Representations

Arnab Das¹, Carlos Franzreb¹, Tim Herzig¹, Philipp Pirlet¹, Tim Polzehl^{1,2}

¹Speech and Language Technology, German Research Center for Artificial Intelligence (DFKI), Germany ²Quality and Usability Lab, Technische Universität Berlin, Germany

{arnab.das, carlos.franzreb, tim.herzig, philipp.pirlet, tim.polzehl}@dfki.de

Abstract

The growing use of speech-based cloud devices and services has heightened the risk of identity theft and misuse of personal information. Speech anonymization techniques help exercise our right to privacy and shield us from falling prey to such malpractices. In this paper, we propose three speech anonymization systems to be submitted to the Voice Privacy Challenge 2024 and describe them in detail. Voice anonymization systems often lack utility for downstream applications, resulting in issues like poor emotion preservation or low intelligibility. This has led to research focused on balancing the privacy-utility trade-off. We propose two methods, that use the KNN-based voice conversion (VC) system as a core anonymization method and show improved intelligibility and emotion preservation. We also propose to employ a vector quantized mutual information-based VC system that learns to distinguish between speaker and content features and alters speaker information during inference time to achieve speaker anonymity. We evaluate these two types of voice conversion systems within the framework of speaker anonymization and analyze the utility-privacy trade-off achieved by each system.

Index Terms: voice anonymization, privacy-utility tradeoff

1. Introduction

In recent times the proliferation of the use of voice-controlled cloud devices or speech-based online services is distinctively noticeable. This can be attributed to several things, but the most significant one is the rapidly expanding field of generative AI research for speech processing and human-computer interaction in general [1]. These speech AI technologies have the potential to significantly transform vital service sectors, such as health-care. Some of the applications could be therapeutic support for mental health patients, automatic language translation, continuous patient monitoring, automated transcription of patient-doctor interactions for future reference, enhanced telemedicine experiences for patients, and many more. Such voice technologies can be extremely beneficial to humans, but they also raise certain privacy concerns. Voice biometrics, like fingerprints, carry behavioral or physiological measurements of individuals [2], which can lead to identity theft by ill-intentioned persons to impersonate a person to acquire wrongful access to services and information. Additionally, a person’s voice can be exploited to create misinformation, disinformation, and hateful content, which can disrupt society. To guarantee privacy and security, protections against such malpractices are crucial.

In this context, voice anonymization can be helpful to provide speaker privacy [3]. It is a technique that modifies the speech signal in such a way that linguistic features remain intact, while the speaker’s characteristics are altered, preventing

the original speaker from being re-identified [3]. The foremost goal of speaker anonymization is to ensure that the speaker cannot be identified purely by their voice thus reducing the chances of voice phishing, impersonation, and other voice-based fraudulent activities. Nowadays, voice anonymization is not only a choice but in some regions, it is mandated by law for data storing to ensure privacy for example the General Data Protection Regulation (GDPR) in Europe [4]. In recent years, generative AI-based voice anonymization research has advanced significantly, and it is capable of producing high-quality anonymized speech signals [1]. However, these deep-learning models employed for speaker anonymization to ensure privacy are often criticized for generating speech signals that are of low naturalness and, intelligibility thus diminishing their utility for downstream applications [5]. This phenomenon, coined as *privacy-utility trade-off* has resulted in a new research direction in the speech privacy community. One challenge is to find appropriate metrics to measure utility [5]. The Voice Privacy Challenge (VPC) initiative [6], founded in 2020, is leading efforts to develop privacy-preserving solutions for voice technologies and also providing necessary metrics for utility quantification [7].

To achieve the objective of speaker anonymization, the baselines provided by VPC mainly follow two major tracks - signal processing techniques and deep learning (DL) based speech synthesis with a mechanism to alter speaker encodings [5], though the results from earlier challenge instances show that the later can produce better results. Signal processing methods include anonymization techniques using McAdams coefficients [8], which adjusts the pole locations generated by linear predictive coding (LPC) analysis of speech input. In this category, another method was proposed in [9], where pitch shifting by timescale modification (TSM) is used for anonymization. These methods don’t require any prior training. The second approach follows a generic scheme of extracting linguistic content feature, speaker embedding and a temporal pitch feature from the input speech signal and reconstructing back anonymized speech after manipulating the speaker embeddings [5]. Following this scheme, [10] proposes a method where the X-vector [11] embedding of the speaker is altered from a pool of speakers to achieve necessary anonymization. [12] uses an extension of ECAPA-TDNN (Time Delay Neural Network) [13] speaker features instead of x-vectors. Alternatively, [5] tries to learn the latent speaker embedding distribution using a generative adversarial network (GAN) [14], and during anonymization generates sufficiently distant pseudo speaker by sampling the learned distribution. These methods show continuous improvement in speaker anonymization when measured objectively however they still lack sufficient intelligibility with high word error rates (WER). Additionally, [15, 16] explore methods involving voice conversion (VC) techniques and show significant improvement

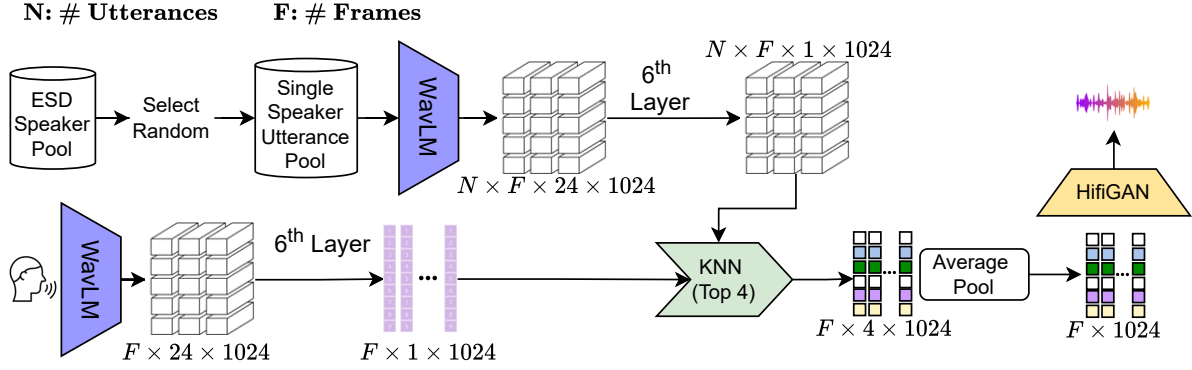


Figure 1: System architecture diagram for KNNS

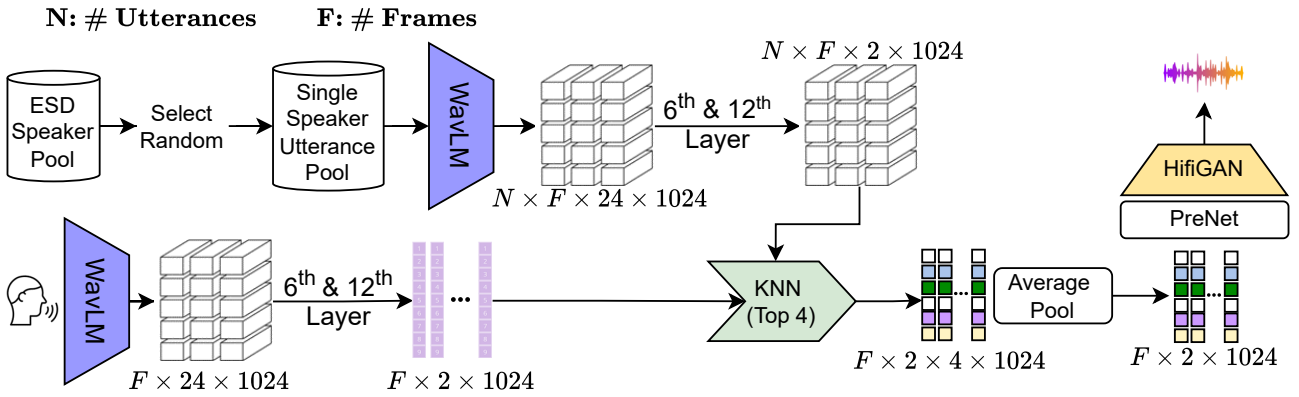


Figure 2: System architecture diagram for KNND

regarding utility measurements.

In line with this, in this paper, we compare two genres of voice conversion techniques to improve upon the utility of anonymized speech. We propose two systems involving knn-vc [17], where the VC component doesn't require any pertaining and built upon speech feature extracted by pre-trained foundation audio models. The third system is built upon VQMIVC [18], which relies on disentangled orthogonal feature decomposition by minimizing mutual information in each pair of features. This method also used vector quantization as a further means of bottleneck for linguistic content features. We discuss these methods in detail in Section 2 and then in Section 4 we present and discuss the results achieved by VPC evaluation scripts.

2. System Description

We propose three anonymization systems for VPC 2024. As described in the VPC 2024 evaluation plan all of the proposed systems are designed to achieve utterance level anonymization.

2.1. Anonymization using single layer knn-vc: KNNS

The core idea of this system is based on knn-vc as proposed in [17]. The system architecture diagram is depicted in Figure 1.

The source utterance is first processed via WavLM Large [19], which outputs speech representation of size $F \times 24 \times 1024$, where F is the number of frames, 1024 is the latent dimension and 24 signifies the output of all transformer layers. WavLM outputs 50 frames from a 1 s long waveform sampled at 16 kHz rate. Then as advised in knn-vc, only the output of the 6th layer is selected, which makes the speech representation of size $F \times 1 \times 1024$, called as *query set*. For anonymization, first, a random target speaker is chosen from all the English speakers present in the emotional speech database (ESD)[20]. Consequently, all the utterances of the chosen target speaker are processed through WavLM, which gives rise to a large feature array of size $N \times F \times 24 \times 1024$, where N is the number of utterances for the chosen random target speaker. Again we chose only the output of the 6th, which makes the feature array of size $N \times F \times 1 \times 1024$. This is called the *matching set*. Then a KNN operation is performed for each frame of the query set to the matching set and the top-4 nearest neighbors are selected, producing a *matched set* of size $F \times 4 \times 1024$. Then it goes through an average pooling operation to average the top 4 neighbors bringing down the matched set size to $F \times 1024$. Finally, a HiFiGAN [21] is applied to reconstruct back the audio signal in the time domain. The HiFiGAN is pre-trained on Librispeech: train-clean-360 [22], ESD [20], CREMA-D [23] and

Table 1: Modules and training corpora for the anonymization system VMC

Module	Description	Output	Training Data
F0 Extractor	Pyworld dio and stonemask Input: Raw waveform	Normalized F0 contour of shape $F \times 1$	-
Speaker Encoder	Convolution Layers + temporal global pool Input: Mel Spectrogram of shape $F \times 80$, with hop length 160 and 80 Mel bands	1×256 speaker feature	LibriSpeech:train-clean-360 ESD RAVDESS CREMA-D
Content Encoder	Convolution Layers + vector quantization Input: Mel Spectrogram of shape $F \times 80$	$F/2 \times 512$ vector quantized content feature	

RAVDESS [24] datasets.

2.2. Anonymization using multi-layer knn-vc: KNND

The working principle of this KNND system is similar to the previously discussed KNNS system in 2.1 and the system diagram is shown in Figure 2. The difference is, that instead of selecting only the output of the 6th layer of WavLM, we select the output of both 6th and 12th layer. In our experiments, we discovered that the 12th layer encodes emotional cues more effectively than the 6th layer. Therefore, we opted to incorporate both layers to enhance emotion preservation. This increases the size of the query set and matching set to $F \times 2 \times 1024$ and $N \times F \times 2 \times 1024$ respectively. Consequently, the matched feature also becomes an array of size $F \times 2 \times 1024$. To accommodate this additional dimension, the HiFiGAN is augmented with a convolutional *PreNet* module, which gets pre-trained jointly with the HiFiGAN before applying to the anonymization pipeline. In both KNNS and KNND, except the HiFiGAN component nothing else is trainable.

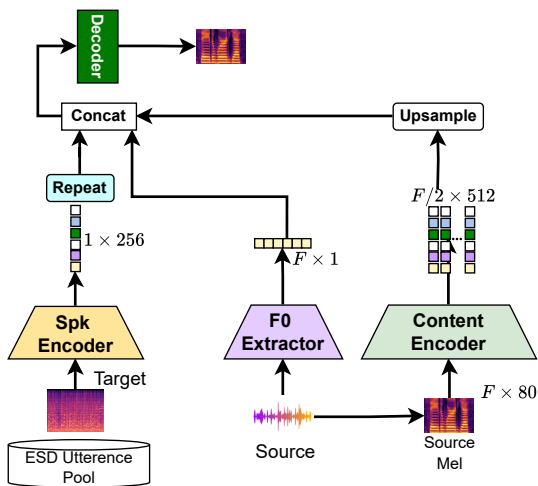


Figure 3: System architecture diagram for VMC

2.3. Anonymization by disentangled representation using vqmivc: VMC

The core of the vqmivc-based (in short VMC) anonymization systems is a VC system proposed in [18] which disentangles speech representations using vector quantization and lowering mutual information in an unsupervised way. The anonymization pipeline for the VMC system is depicted in Figure 3. The core VC framework has four major components, an F0 extractor, a speaker encoder, a content encoder, and a decoder.

Step I - Training: During training the F0 extractor receives the raw source waveform and uses the *pyworld-dio algorithm* to extract temporal F0 features of shape $F \times 1$, where F is the number of frames. The speaker encoder takes input from a Mel spectrogram of the source utterance of shape $F \times 80$ and produces a global speaker feature vector of shape 1×256 . The content encoder also receives the same Mel spectrogram and produces content feature representation of size $F/2 \times 512$. The content encoder comprises convolution layers and a codebook to quantize the content features. The disentanglement of F0 features, speaker features, and content features is achieved by minimizing parameterized mutual information as described in [18]. Afterward, the content feature is upsampled to achieve the shape $F \times 512$ and the global speaker feature is repeated in temporal dimension to achieve the shape of $F \times 256$. Consequently, the three speech features are concatenated and fed to a decoder which outputs the original Mel spectrogram. The whole VC framework is trained unsupervised by optimizing reconstruction loss. We have kept the architecture of each module the same as recommended in [18]. Implementation details of individual components are outlined in Table 1. Finally, a HiFiGAN is applied to reconstruct back the waveform from the Mel spectrogram. Similar to KNNS and KNND systems, the HiFiGAN is pre-trained on Librispeech: train-clean-360, ESD, CREMA-D, and RAVDESS datasets.

Step II - Inference: During the anonymization, a random target utterance is selected from the ESD utterance pool, and the speaker encoder is fed with the Mel spectrogram of the target utterance whereas the other two components F0 extractor and the content encoder receive the source utterance. This produces the output utterance in the target speaker's voice thus effectively archiving anonymization.

Table 2: Objective evaluation results for privacy metric - EER using the semi-informed scenario. Among the proposed systems, the best scores are in bold, and the second-best scores are underlined per gender.

Systems	ASV EER \uparrow			
	Librispeech			
	Dev		Test	
	Female	Male	Female	Male
McAdams - B2 [7]	12.91	2.05	7.48	1.56
GAN - B3 [7]	28.43	22.04	27.92	26.72
KNNS	12.64	4.19	9.51	4.45
KNND	17.07	8.54	<u>10.77</u>	<u>9.36</u>
VMC	19.462	8.54	16.79	11.8

3. Experiment Setup

3.1. Datasets

For training the HiFiGAN modules in all three systems and the VQMIVC voice converted for the VMC systems we used in total of four datasets - Librispeech: train-clean-360 [22], ESD [20], CREMA-D [23] and RAVDESS [24]. The train-clean-360 subset contains more than 960 hours of clean speech data from 921 speakers (439 female and 482 male) with an average of 25 minutes of speech per speaker. The ESD database includes parallel utterances conveying 5 emotion categories (neutral, happy, angry, sad, and surprised), spoken by 10 native English (gender-balanced) speakers. CREMA-D dataset contains > 7000 clips from 91 actors (48 male and 43 female) covering 6 different emotions - Anger, Disgust, Fear, Happy, Neutral, and Sad. In RAVDESS, 24 professional actors (12 females and 12 males) uttered parallel emotional contentment depicting 7 different emotions - calm, happiness, sadness, anger, fear, surprise, and disgust. The evaluation of the proposed systems is performed on the dev-clean and the test-clean subsets from the Librispeech corpus and the IEMOCAP [25] dataset is used as prescribed by the VPC challenge.

3.2. Objective metrics

The anonymization systems are evaluated using the models suggested by the VPC challenge. To objectively measure the privacy benefits, an attack using automatic speaker verification (ASV) is applied and an equal error rate (EER) score is computed. The attack uses both a semi-informed scenario alternative to an ignorant one. In the ignorant case, the ASV model is trained on real utterances and applied to anonymized utterances. Alternatively, in the semi-informed scenario, the ASV model is trained on the anonymized utterances from the training subset and then applied to the utterances from dev-clean and test-clean subsets. Both of these scenarios consider the anonymization system as a black box. Evidence shows that a semi-informed attack is stronger than an ignorant one.

To measure the utility of the anonymized utterances, automatic speech recognition (ASR) and speech emotion recognition (SER) models are applied to judge the intelligibility and amount of emotion preservation respectively. For intelligibility, word error rate (WER), and for emotion preservation un-weighted average recall (UAR) is reported.

Table 3: Objective evaluation results for utility metrics - WER for ASR and mean UAR (5 folds) for SER. Among the proposed systems, the best scores are in bold, and the second-best scores are underlined.

Systems	SER UAR [%] \uparrow		ASR WER [%] \downarrow	
	IEMOCAP		Librispeech	
	Dev	Test	Dev	Test
Original [7]	69.08	71.06	1.8	1.85
McAdams - B2 [7]	55.61	53.49	10.44	9.95
GAN - B3 [7]	38.09	37.57	4.29	4.35
KNNS	46.94	47.69	<u>2.6</u>	<u>2.52</u>
KNND	<u>43.2</u>	<u>45.54</u>	2.47	2.49
VMC	35.14	33.92	20.75	18.25

4. Results & Discussion

The results for objective evaluation to measure privacy strength are summarized in Table 2. We compare our proposed systems with the signal processing dependent McAdams (system ID - **B2**) baseline and DL model dependent GAN-based anonymization (system ID - **B3**) baseline system as reported in [7]. The attack scenario is semi-informed as the ASV system is trained on anonymized data. *B2* archives EER scores of 12.91 and 2.05 in the dev subset for female and male trail utterances respectively. On the test subset, the EER scores are even lower, 7.48 and 1.56 respectively for female and male trials. GAN-based *B3* baseline archives much higher EER scores indicating stronger anonymization. For *B3*, EER scores in the dev subset are 28.43 and 22.04, whereas for the test subset, the scores are 27.92 and 26.72 for female and male trials. Among the 3 proposed systems, VMC achieves the highest EER scores for both the subsets and KNND EER scores are the second best. For the dev subset, VMC archives EER scores of 19.462 and 8.54, closely followed by the KNND system, which archives 17.07 and 8.54 EER scores for female and male trail utterances. For the test subset, EER scores for the VMC system are 16.79 and 11.8 and for the KNND system the scores are 10.77 and 9.36. Comparatively, EER scores for the *KNNS* system are a bit low. It achieves EER scores of 12.64 and 4.19 on the dev subset and for the test subset, the scores are 9.51 and 4.45. So in terms of anonymization strength, all the proposed systems beat the *B2* baseline but are somewhat inferior to the *B3* baseline. Noticeably, for all the systems including the baseline EER score for Female trial utterances are higher than that of males, which indicates it's easier to anonymize female speakers than male speakers.

The results of the objective utility evaluation are presented in Table 3. In terms of UAR for the SER task, the scores are 69.08% and 71.06% for the dev and test subsets of IEMOCAP data respectively when tested on original utterances. However, the *B2* baseline archives UAR scores of 55.61 and 53.49 whereas the scores achieved by the *B3* baseline are way lower, 38.09 and 37.57 for the dev and test subset respectively as reported in [7]. Among the proposed systems, KNNS achieves the highest UAR scores for the SER task closely followed by the KNND system. The KNNS system archives 46.94% and 47.69% on the dev and test subset respectively whereas the KNND system achieves 43.2% and 45.54%. The emotion preservation capability for the VMC system is comparatively lower as it achieves 35.14% and 33.92% on the IEMOCAP dev and test subsets respectively.

In terms of intelligibility evaluation, the original data get

WER scores of 1.8% and 1.85% on Librispeech dev and test subsets respectively. For the *B2* baseline the scores are 10.44% and 9.95% however the *B3* baseline achieves better intelligibility as the WER scores are 4.29% and 4.35%. This shows that even though the *B3* baseline is not as good as *B2* in terms of emotion preservation for intelligibility of anonymized speech *B3* is much better than *B2*. Among the proposed systems, both KNND and KNNS systems significantly beat both the baseline systems as the WER scores are very low and almost analogous to the original data. The best WER scores are achieved by the KNND system, 2.47% and 2.49% respectively on the dev and test subset, whereas KNNS achieves 2.6% and 2.52%. Similar to the SER task, VMC is lacking on the ASR task as well, as it achieves high WER scores of 20.75% and 18.25%.

The privacy and utility results presented in Table 2 and Table 3 clearly show that we get higher privacy at the cost of lower utility of the anonymized speech. For the GAN-based *B2* baseline system the privacy scores are high but the utility scores are low, especially for the emotion preservation task. Our proposed KNND system achieves a good balance of privacy and utility dimensions, as evidenced by significantly lower WER scores than all other systems, higher emotion preservation compared to the *B3* baseline, and higher privacy compared to the *B2* baseline. The proposed KNNS is also capable of producing high utility scores for the anonymized data as it beats the *B3* baseline in both SER and ASR tasks however in terms of speaker privacy the performance is a bit lower than the KNND system. The VMC system, which archives anonymization by disentangling content and speaker features, can archive high privacy for anonymized speech compared to KNNS and KNND but has poor utility.

5. Conclusion

In this paper, we proposed three anonymization systems by VC, two of them are based on learnable parameters less knn-vc, and the last one is based on VQMIVC with learns to disentangle speaker and content features. Our proposed knn-vc based systems use pre-trained WavLM features as speech encoding and the results show that they achieve a very good balance between the privacy and utility of the anonymized speech, especially the KNND system. Both the KNNS and KNND systems are capable of producing highly intelligible anonymized speech, as their WER scores are marginally shy of the WER scores of the original speech and significantly better than the baseline systems. The third system shows that the feature disentanglement can provide higher privacy against an attacker but the anonymized speech lacks utility as both emotion preservation and intelligibility are poor. We will take this up as future work to improve the utility of such an end-to-end feature disentanglement-based anonymization system without losing its capability of providing good privacy.

6. Acknowledgements

This research has been partly funded by the Federal Ministry of Education and Research Germany (BMBF 16KISA007, project Medinym) and partly by the Volkswagen Foundation.

7. References

- [1] T. Feng, R. Hebban, N. Mehlman, X. Shi, A. Kommineni, S. Narayanan *et al.*, “A review of speech-centric trustworthy machine learning: Privacy, safety, and fairness,” *APSIPA Transactions on Signal and Information Processing*, vol. 12, no. 3, 2023.
- [2] N. Singh, A. Agrawal, and R. Khan, “Voice biometric: A technology for voice based authentication,” *Advanced Science, Engineering and Medicine*, vol. 10, no. 7-8, pp. 754–759, 2018.
- [3] M. Maoche, B. M. L. Srivastava, N. Vauquier, A. Bellet, M. Tommasi, and E. Vincent, “A comparative study of speech anonymization metrics,” in *INTERSPEECH 2020*, 2020.
- [4] I. Siegert, V. S. Varod, N. Carmi, and P. Kamocki, “Personal data protection and academia: Gdpr issues and multi-modal data-collections,” *Online Journal of Applied Knowledge Management (OJAKM)*, vol. 8, no. 1, pp. 16–31, 2020.
- [5] 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)*. IEEE, 2023, pp. 912–919.
- [6] N. Tomashenko, B. M. L. Srivastava, X. Wang, E. Vincent, A. Nautsch, J. Yamagishi, N. Evans, J. Patino, J.-F. Bonastre, P.-G. Noé *et al.*, “Introducing the voiceprivacy initiative,” *arXiv preprint arXiv:2005.01387*, 2020.
- [7] 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.
- [8] J. Patino, N. Tomashenko, M. Todisco, A. Nautsch, and N. Evans, “Speaker anonymisation using the mcadams coefficient,” *arXiv preprint arXiv:2011.01130*, 2020.
- [9] C. O. Mawalim, S. Okada, and M. Unoki, “Speaker anonymization by pitch shifting based on time-scale modification,” in *2nd Symposium on Security and Privacy in Speech Communication*. ISCA, 2022.
- [10] F. Fang, X. Wang, J. Yamagishi, I. Echizen, M. Todisco, N. Evans, and J.-F. Bonastre, “Speaker anonymization using x-vector and neural waveform models,” *arXiv preprint arXiv:1905.13561*, 2019.
- [11] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust dnn embeddings for speaker recognition,” in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [12] R. Khamsehashari, Y. Sinha, J. Hintz, S. Ghosh, T. Polzehl, C. Franzreb, S. Stober, and I. Siegert, “Voice privacy-leveraging multi-scale blocks with ecapa-tdnn se-res2next extension for speaker anonymization,” in *Proc. 2nd Symposium on Security and Privacy in Speech Communication*, 2022, pp. 43–48.
- [13] B. Desplanques, J. Thienpondt, and K. Demuynck, “Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” *arXiv preprint arXiv:2005.07143*, 2020.
- [14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [15] H.-P. Chang, I.-C. Yoo, C. Jeong, and D. Yook, “Zero-shot unseen speaker anonymization via voice conversion,” *IEEE Access*, vol. 10, pp. 130 190–130 199, 2022.
- [16] M. Costante, M. Matassoni, and A. Brutti, “Using seq2seq voice conversion with pre-trained representations for audio anonymization: experimental insights,” in *2022 IEEE International Smart Cities Conference (ISC2)*. IEEE, 2022, pp. 1–7.
- [17] M. Baas, B. van Niekerk, and H. Kamper, “Voice conversion with just nearest neighbors,” *arXiv preprint arXiv:2305.18975*, 2023.
- [18] D. Wang, L. Deng, Y. T. Yeung, X. Chen, X. Liu, and H. Meng, “Vqmivc: Vector quantization and mutual information-based unsupervised speech representation disentanglement for one-shot voice conversion,” *arXiv preprint arXiv:2106.10132*, 2021.
- [19] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.

- [20] K. Zhou, B. Sisman, R. Liu, and H. Li, "Emotional voice conversion: Theory, databases and esd," *Speech Communication*, vol. 137, pp. 1–18, 2022.
- [21] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," *Advances in neural information processing systems*, vol. 33, pp. 17 022–17 033, 2020.
- [22] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [23] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "Crema-d: Crowd-sourced emotional multimodal actors dataset," *IEEE transactions on affective computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [24] S. R. Livingstone and F. A. Russo, "The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english," *PloS one*, vol. 13, no. 5, p. e0196391, 2018.
- [25] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, pp. 335–359, 2008.