

# HLTCOE JHU Submission to the Voice Privacy Challenge 2024

Henry Li Xinyuan, Zexin Cai, Ashi Garg, Kevin Duh, Leibny Paola García-Perera, Sanjeev Khudanpur, Nicholas Andrews, Matthew Wiesner

Johns Hopkins University, United States

xli257@jhu.edu

## Abstract

We present a number of systems for the Voice Privacy Challenge, including voice conversion based systems such as the kNN-VC method and the WavLM voice Conversion method, and text-to-speech (TTS) based systems including Whisper-VITS. We found that while voice conversion systems better preserve emotional content, they struggle to conceal speaker identity in semi-white-box attack scenarios; conversely, TTS methods perform better at anonymization and worse at emotion preservation. Finally, we propose a random admixture system which seeks to balance out the strengths and weaknesses of the two category of systems, achieving a strong EER of over 40% while maintaining UAR at a respectable 47%.

**Index Terms:** Voice Conversion, Text-to-Speech, Voice Anonymization

## 1. Introduction

Technology has enabled unprecedented means for extracting personally identifiable information from speech. New legislation that recognizes this reality including the Illinois biometric privacy act [1], California’s consumer privacy act of 2018, South Africa’s Protection of personal information act (POPIA), and most notably, the EU’s General Data Protection Regulation (GDPR), greatly limit the extent to which personally identifiable information, including speech can be stored or processed. Compliance with these laws, and protection of users’ private information requires means of anonymizing a user’s voice without corrupting the data to such an extent that it no longer has any linguistic or paralinguistic content. The Voice Privacy Challenge is one initiative driving these efforts.

In the Voice Privacy Challenge voice anonymization is modeled as a user-attacker game. A *user* anonymizes their speech with the goal of preserving linguistic and emotional content, while an *attacker* attempts to identify speech from any remaining signal in the anonymized utterances including the remaining linguistic and emotional content. The challenge specifies that the Voice Anonymization system is a function  $f$  that operates on the utterance level.

To measure these two competing objectives, challenge submissions are evaluated using utility and privacy metrics. Utility refers to the extent to which the spoken and emotional is preserved following anonymization. These metrics are measured using a pre-trained systems: we note that this approach makes it possible to optimize directly on these metrics. Content preservation is measured by Word Error Rate (WER) in a traditional Automatic Speech Recognition (ASR) task setup. Emotion is measured by Unweighted (across the different emotion labels) Average Recall (UAR). Privacy is measured using an ECAPA-TDNN-based [2] Automatic Speaker Verification (ASV) sys-

tem. At training time, a classifier is attached to the ECAPA-TDNN model and trained jointly on speaker labels; at inference time, ECAPA-TDNN representations are extracted for a set of “enrollment utterances” — utterances known to have been produced by a certain speaker, and for “test utterances” - utterances whose speaker the attacker endeavors to determine. A distance metric (cosine distance) is computed for each test utterance relative to the pool of enrollment utterances. A distance threshold, where accuracy equals recall if all utterances whose distance to the enrollment pool is less than the threshold are taken to belong to enrollment speaker, is identified. The resulting error rate is known as Equal Error Rate (EER). The theoretical upper limit of EER is 50%.

Anonymization systems tend to excel at either utility preservation or anonymization, but not both. To address this short-coming, we propose two different anonymization methods that excel on each individual task and a simple, but effective method to combine these approaches that enables a reasonable tradeoff between the competing objectives. First we describe a simple voice-conversion approach that focuses on preserving linguistic and emotional content, but not anonymization under the strong attack model in the Voice Privacy Challenge. Inspired by the anonymization approach by Meyer et al. [3], we also adopt a cascaded ASR-TTS approach for anonymization. However, the original approach replicates the prosodic features of the original speech, which can lead to speaker identity leakage in this year’s challenge, where attackers have access to anonymized speech for training speaker verification models. To mitigate this, we use a multi-speaker TTS system to eliminate the speaker footprint. This cascaded approach excels in anonymization but degrades the utility of the speech signal. Finally, we describe our method for combining both approaches.

## 2. Method

### 2.1. Our Baseline: kNN-VC

Our baseline system extends earlier work by Baas et al. [4]. The method, as with our subsequent adaptations, operate at WavLM-feature level [5]. The original kNN-VC system consists of a simple yet elegant idea: first, both the source utterance and a number of utterances from the target speaker (totalling at least 5 minutes in length) are converted into the WavLM feature space; next, k-nearest neighbour regression is performed on each frame in the WavLM-feature representation of an utterance with respect to the set of all WavLM feature frames in the target speaker’s utterance pool. The average of the k-nearest neighbour output is used to synthesize the target utterance using a HiFi-GAN vocoder [6] trained to synthesize speech using WavLM feature vectors. The vocoder we used in our system was trained on the LibriSpeech train-clean-100 subset, match-

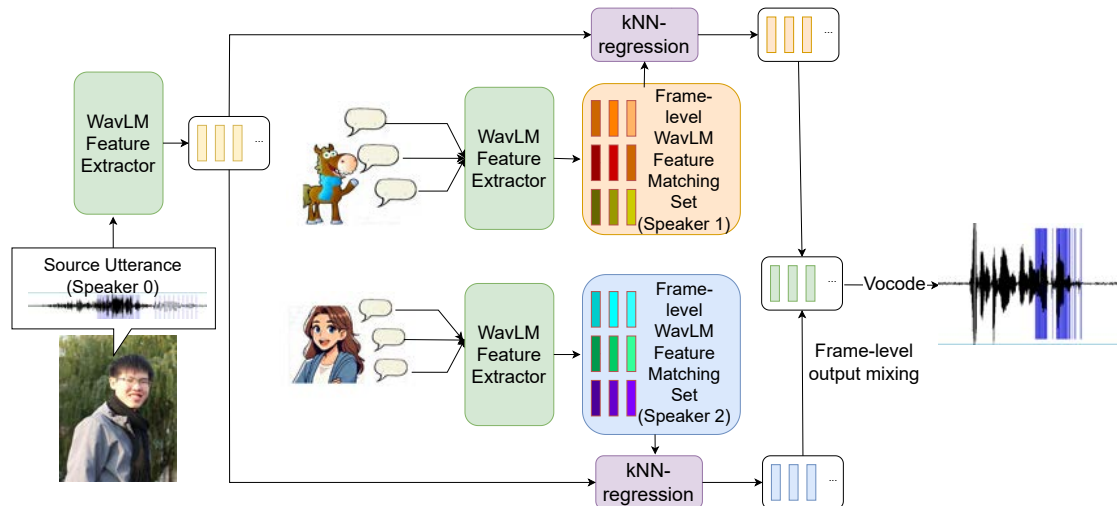


Figure 1: Schematic of our adapted kNN-VC system.

ing the one used by the original authors of kNN-VC. We picked  $k = 4$  during our k-nearest neighbour regression, matching the setup in the original paper.

The following modifications to the original kNN-VC system (ID = 1 in Table 1) were shown to perform marginally superior at anonymization: choosing random target speakers; applying length variation (ID = 2); applying an additive noise to the WavLM features pre-conversion (ID = 3). These improvements come at the cost of emotion or content preservation.

The original kNN-VC system targets the voice characteristics of a single speaker, which may run contrary to the overarching goals of privacy protection. In order to address this problem, we revised the kNN-VC system so that an arbitrary number of matching sets, each originating from a different speaker/pseudo-speaker, is prepared. At inference time, the kNN outputs from each matching set is pooled together using a weighted average. The schematic of our revised kNN-VC system is shown in figure 1.

## 2.2. WavLM Conversion

The kNN-VC system produces converted utterances that, qualitatively, have remarkably similar vocal characteristics to those produced by the target speaker. It is strong at preserving both the content and the emotion of the original utterance, outperforming all of the previous submissions. We do, however, note the relative weakness of the kNN-VC system on the privacy objective compared to cascaded TTS-based systems (such as [7]). The simplicity of this method also hinders any efforts to address this weakness. As such, we propose a transformer-based WavLM feature conversion model, in the hopes that a neural method would offer more flexibility to allow us to tackle this weakness.

One obstacle to training voice conversion models is the lack of parallel data from different speakers. Fortunately, our baseline kNN-VC system provides us with a virtually unlimited synthetic pool of such training data. As such, we propose the following definitions:

1.  $w$ : WavLM feature extractor
2.  $k$ : kNN-VC feature converter.  $k(w(u), a)$  converts the WavLM features of utterance  $u$  into that of target speaker

$a$  using the kNN-VC system.

Given these definitions, we proceed to train our WavLM feature conversion system to predict  $k(w(u), a)$  from  $w(u)$ .

The base model architecture was adapted from FastSpeech2 [8], a non auto-regressive, encoder-decoder TTS system. We adapted an implementation of FastSpeech2 from Chien et al. [9], modifying the model task from TTS to Voice Conversion and changing the representation that the model operates on from Mel Spectrogram to WavLM features. Our model uses 4 encoder layers and 6 decoder layers, each with a hidden dimension of 512.

In its most basic form, the model learns to faithfully replicate the output and therefore inheriting the weakness at anonymization of the kNN-VC system. Changes to the training procedure are necessary so as to improve in that regard.

Some of the additions we made to the base WavLM conversion model include:

1. Joint training with target reconstruction objective (such that some  $u$  belong to the target speaker ( $u_a$ ))
2. Joint Adversarial training with speaker ID task. A speaker-ID system is jointly trained with the WavLM conversion model. The WavLM conversion model is optimized to hinder the speaker-ID system from producing the correct prediction.
3. Discretized and aligned objective

We will proceed to describe each of these additions in details:

### 2.2.1. Joint training with target reconstruction objective

In every batch, a percentage of  $(w(u), k(w(u)))$  pairs are replaced by  $(w(u_a), w(u_a))$ , where  $u_a$  denotes an actual utterance by the target speaker. We hope that, by inducing the model to learn to reconstruct large chunks of the target speaker’s speech, it would pick up on sequence-level characteristics of the target speaker’s speech patterns and reproduce them at inference time during voice conversion.

### 2.2.2. Joint adversarial training

An ECAPA-TDNN-based Speaker-ID model, similar to the one that is pre-trained prior to being used for ASV evaluation in

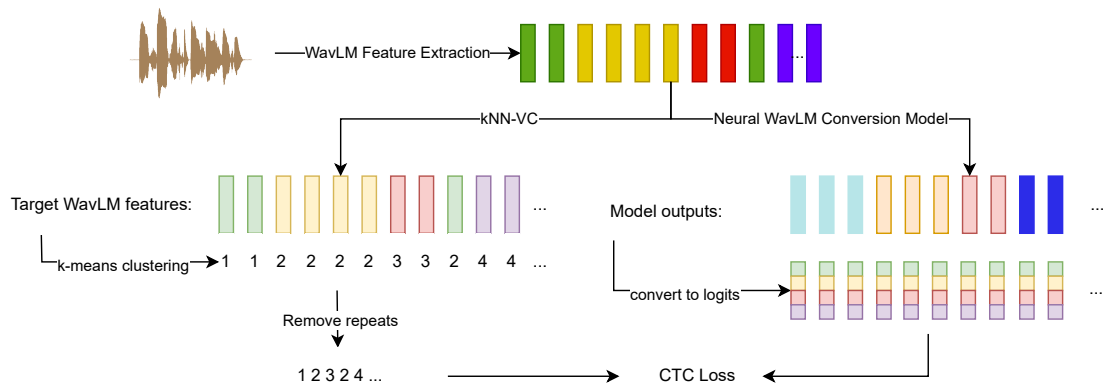


Figure 2: Schematic of our WavLM conversion system with  $k$ -means discrete loss. The training targets (on the left hand side) are discretized using  $k$ -means clustering. The resulting token sequence is used as the golden target labels during CTC loss calculation.

the challenge, can be trained to work with WavLM features as well, achieving up to 99.6% accuracy on the train-clean-360 split of LibriSpeech [10]. We leverage this fact to perform Joint Adversarial Training between voice conversion and speaker ID. Specifically, the speaker ID system and the voice conversion system are trained jointly. For each batch of samples, the speaker ID system gets its weights updated so that it learns to better classify the source speakers using the output of the voice conversion system; the voice conversion system, in addition to learning to voice convert, also learns to prevent the speaker ID system from being able to correctly classify the source speaker.

### 2.2.3. Discretized and aligned objective

The original FastSpeech2 implementation used a combination of L1 and L2 losses against the training target - utterances converted using the kNN-VC system. As the loss operates on the frame-level, it encourages the model to preserve the sequence-level information of the source utterance such as speed or accent: information that leads to strong emotion preservation at the expense of weaker anonymization. In order to allow the model to warp an utterance in the time domain, we design the following training objective: first, we discretize the set of target speaker WavLM feature vectors using  $k$ -means clustering. Using the resulting  $k$ -means model, we discretize each kNN-VC converted source utterance into sequences of discrete tokens. We then collapse any repeated tokens in each discrete label sequence. At training time, cosine distance between each frame that the WavLM conversion model generates and each  $k$ -means centroid is used as a logit value, which is in turn used to compute a likelihood vector for each centroid at each frame. A CTC-loss [11] is then computed using the frame-level likelihood vectors against the collapsed target-side discrete unit sequence. Figure 2 illustrates our WavLM conversion system with the discretized objective in action.

## 2.3. Cascaded Anonymization

Although voice conversion systems can effectively alter the acoustic characteristics related to the speech production organs of source speakers, the prosodic characteristics, which reflect their acquired speaking habits and styles, remain unchanged. We hypothesize that such prosodic characteristics are benefi-

cial for preserving emotions, while also serving as patterns to identify the corresponding speaker. To address this, we cascade automatic speech recognition (ASR) and text-to-speech (TTS) models to enhance anonymization performance by altering the speaking style of the source speech.

The anonymization process of our cascaded system is illustrated in Figure 3. We first obtain the transcript of the source utterance using an open-sourced ASR system. Then, we apply a TTS system to generate the corresponding anonymized utterance with a randomized voice. Specifically, we used the ‘medium-en’ model from Whisper<sup>1</sup> [12] as our ASR system to obtain the transcript. The synthesis system is an open-source multi-speaker TTS model<sup>2</sup>, VITS [13], trained on the LibriTTS dataset. For anonymization purposes, we randomly selected a speaker’s voice from a pool of 904 speakers to synthesize each utterance.

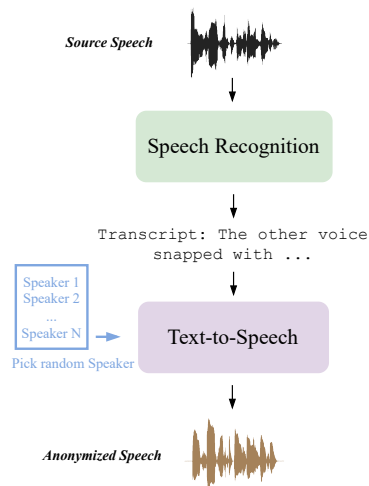


Figure 3: Cascaded ASR-TTS Anonymization Process

<sup>1</sup><https://github.com/openai/whisper>  
<sup>2</sup>[https://huggingface.co/datasets/rhasspy/piper-checkpoints/blob/main/en/en\\_US/libritts\\_r/medium](https://huggingface.co/datasets/rhasspy/piper-checkpoints/blob/main/en/en_US/libritts_r/medium)

## 2.4. Random Admixture

As discussed earlier, our different systems exhibited different strengths and weaknesses: some excel at anonymization while others appear to do better at emotion preservation. Drawing inspiration from data poisoning attacks against neural networks [14], which found that inserting poor quality or incorrect data into the training set can have a disproportional impact on model performance, we experiment with randomly selecting from one of our many anonymization systems during training at test time. We highlight the random admixture system where the two source pools we drew from are the Whisper-VITS TTS system and the original kNN-VC system. We note that while the emotion preservation performance is a linear extrapolation of that of the two source systems, the trained ASV system performed worse than the linear extrapolation of the EER of these two systems, as studies on data poisoning attacks would predict.

# 3. Experiments and Results

## 3.1. Datasets

Subsets from the LibriSpeech [10] and IEMOCAP [15] corpora are used as development and evaluation data in the challenge. Detailed information can be found in the data description section of the challenge’s evaluation plan [16]. Specifically, there are 10 subsets for the evaluation process. The subsets `libri-dev-asr` and `libri-test-asr` are used for ASR evaluation. The subsets `libri-dev-enrolls`, `libri-dev-trials-f`, `libri-dev-trials-m`, `libri-test-enrolls`, `libri-test-trials-f`, and `libri-test-trials-m` are used for evaluating anonymization (speaker verification) performance. The subset `libri-train-clean-360` is used for training the speaker verification system after anonymization. For emotion preservation performance, `IEMOCAP-dev` and `IEMOCAP-test` are used. For our voice conversion systems, including our submitted kNN-VC system, we drew our target speaker(s) from the VoxCeleb [17] dataset.

## 3.2. Evaluation Metrics

### 3.2.1. Privacy

The main metric used for privacy evaluation is the equal error rate (EER). EER is calculated based on similarity scores from each pair of utterances in the evaluation set. Such a pair is also known as a trial. The threshold  $\theta$  is denoted as the decision boundary between same-speaker and different-speaker, while  $P_{fa}(\theta)$  and  $P_{miss}(\theta)$  represent the false alarm and miss rates at threshold  $\theta$ , respectively. The EER corresponds to the threshold  $\theta_{EER}$  where  $P_{fa}(\theta_{EER}) = P_{miss}(\theta_{EER})$ . A lower EER signifies a higher speaker re-identification risk. Therefore, in the context of voice privacy, a higher EER indicates better privacy.

### 3.2.2. Utility

Two metrics are used to assess the preservation of specific acoustic attributes from the original speech. Word error rate (WER) is adopted to evaluate the ability of an anonymization system to preserve linguistic content. For emotion state preservation, unweighted average recall (UAR) is used for evaluation. Utility performances are evaluated using a pre-trained ASR system and a speech emotion recognition (SER) system, respectively. A lower WER denotes better preservation of linguistic content, while a higher UAR indicates better preservation of emotion states.

## 3.3. Results

The anonymization results of our systems are shown in Table 1. As indicated in the table, the kNN-VC system achieves an average EER of 7.95%, which suggests a significant exposure of the original speaker’s traits. Nevertheless, the kNN-VC system performs well in preserving emotional state and linguistic content, with an average UAR of 56.7% and an average WER of 3.16%, respectively. The model with speed perturbation achieves similar scores to the vanilla kNN-VC system: we found that applying a random factor between 0.8 and 1.2 to the original audio gives us marginal EER improvements without hurting UAR. Applying additive random noise drawn from a uniform distribution of vectors with infinity norm less than 32 on each of the input WavLM feature frames prior to applying kNN-VC system led to greatly improved EER at the expense of reduced UAR and WER.

The improvements we made on top of the kNN-VC system, including the neural WavLM feature conversion system, the speaker adversarial training technique, and the discrete loss, were all shown to moderately improve the anonymization performance of the system, with our WavLM Conversion system combining all features achieving an average EER of 13.88%. However, these improvements come at the cost of greatly reducing the systems’ performance on utility metrics.

On the other hand, the cascaded anonymization approach, `whisper-VITS`, achieves an average EER of 48.25%, which is close to 50%, demonstrating its ability to conceal the original speaker’s identity. Additionally, the WER utility score is 3.75%, which is slightly lower than that of the VC approaches. Note that the WER of the transcribed text by the `whisper` model is 3.38% on the `libri-dev-asr` set and 3.29% on the `libri-test-asr` set. However, the emotion preservation performance of this approach is poor, with the system achieving an average UAR of 30.35% on the IEMOCAP evaluation sets.

By randomly mixing the kNN-VC system with the cascaded anonymization system during voice anonymization, we were able to produce systems whose performance sit in between those of the two systems (as shown in Figure 4), satisfying different needs for anonymization versus emotion preservation. At 50% admixture, the EER of the system approaches 50% while the UAR still stands at a respectable 44.77%, stronger than any existing TTS-based anonymization system.

## 4. Limitations and Future Work

While our Random Admixture system achieved strong results in the context of the Voice Privacy Challenge setup, where the adversary uses a single model trained on anonymized data, we note that an informed adversary can effectively mitigate the impact of data-poisoning attacks at the cost of some performance degradation [18]. We call for future studies to identify any possible equilibrium between the anonymizer and the adversary in light of admixture systems, as well as more generally, the potential role of adversarial attacks within the context of voice privacy.

We note that our TTS system does not support any type of controlled generation, which severely limits its capacity for preserving para-linguistic features such as emotion. We would like to explore investigate the effectiveness of that line of methods in our future work.

Table 1: Privacy and Utility Performance of Various Anonymization Approaches  
(Darker Color Indicates Better Performance)

ID	System	Privacy - EER (%) $\uparrow$					Utility - UAR mean (%) $\uparrow$			Utility - WER (%) $\downarrow$		
		libri-dev-f	libri-dev-m	libri-test-f	libri-test-m	avg.	IEMOCAP-dev	IEMOCAP-test	avg.	libri-dev	libri-test	avg.
0	origin	10.511	0.931	8.761	0.418	5.16	69.0796	71.0618	70.07	1.807	1.844	1.83
1*	kNN-VC	11.789	5.141	9.307	5.570	7.95	56.7330	56.6740	56.70	3.275	3.048	3.16
2	kNN-VC + len variation	11.192	5.125	10.218	5.793	8.08	56.9488	55.638	56.29	3.28	3.387	3.33
3	kNN-VC+ len var + noise-in	24.681	18.624	19.891	19.115	20.58	44.1260	42.3846	43.26	11.993	10.008	11.00
4*	whisper-VITS	47.584	49.233	47.445	48.750	48.25	30.1074	30.5932	30.35	3.743	3.755	3.75
1 + 4*	Admixture ( $p = 0.2$ )	26.003	16.155	20.776	24.722	21.91	51.2840	52.1324	51.71	3.300	3.290	3.31
1 + 4*	Admixture ( $p = 0.325$ )	34.518	32.918	34.532	33.676	33.91	49.3398	48.7304	49.04	3.514	3.336	3.43
1 + 4*	Admixture ( $p = 0.4$ )	41.192	40.660	42.182	39.225	40.81	47.0784	47.1046	47.09	3.454	3.199	3.33
5	WavLM Conv (base)	13.622	6.987	9.307	4.231	8.54	55.5458	53.9522	54.75	3.044	2.982	3.01
6	WavLM Conv + Adv Spk Loss	17.472	9.005	12.773	7.164	11.60	50.7706	50.4628	50.62	4.442	4.015	4.23
7	WavLM Conv + Discrete Loss	18.041	12.268	13.716	10.913	13.73	44.5292	42.5980	43.56	10.313	10.014	10.16
8	WavLM Conv + Adv + Discrete Loss	19.308	11.645	13.870	10.690	13.88	44.0936	42.9102	43.50	10.811	10.850	10.83

\* marks submitted systems

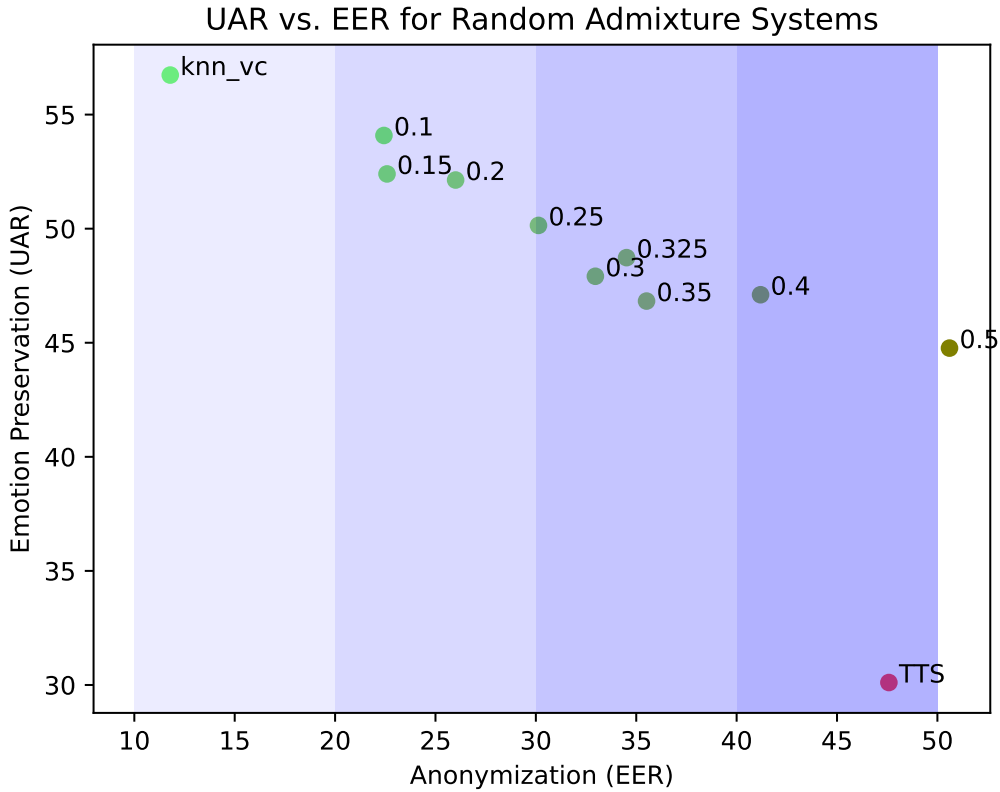


Figure 4: Scatter plot of various results from the Random Admixture system. Results from the two source systems, cascaded TTS and kNN-VC, are included. Each point is labeled and color-coded with the percentage of the admixture which was drawn from the cascaded TTS system.

## 5. Conclusion

We experimented with two main approaches to voice anonymization: voice conversion and cascaded ASR-TTS. Our voice conversion systems generally performed strongly in emotion preservation, while our cascaded anonymization systems excel at anonymization. We were able to freely adjust the trade-off between emotion preservation and anonymization by performing a random admixture of these two systems. We call for future work to investigate if it would be possible to better the

UAR-EER tradeoff curve achieved by random admixture.

## 6. Acknowledgements

This work was supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the ARTS Program under contract D2023-2308110001. The views and conclusions contained herein are those of the authors and should not be interpreted as

necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

## 7. References

- [1] Illinois General Assembly, “Act 14 - Biometric Privacy Act,” <https://www.ilga.gov/legislation/ilcs/ilcs3.asp?ActID=3004&ChapterID=57>, 2008.
- [2] B. Desplanques, J. Thienpondt, and K. Demuynck, “ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification,” in *Interspeech 2020*, pp. 3830–3834.
- [3] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2023, pp. 1–5.
- [4] M. Baas, B. van Niekerk, and H. Kamper, “Voice Conversion With Just Nearest Neighbors,” in *Interspeech 2023*, pp. 2053–2057.
- [5] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, “WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, p. 1505–1518, 2022.
- [6] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity Speech Synthesis,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 17 022–17 033.
- [7] U. E. Gaznepoglu, A. Leschanowsky, and N. Peters, “VoicePrivacy 2022 System Description: Speaker Anonymization with Feature-Matched F0 Trajectories,” in *2nd Symposium on Security and Privacy in Speech Communication*, 2022.
- [8] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, “FastSpeech 2: Fast and High-Quality End-to-End Text to Speech,” in *International Conference on Learning Representations*, 2021.
- [9] C.-M. Chien, J.-H. Lin, C.-y. Huang, P.-c. Hsu, and H.-y. Lee, “Investigating on Incorporating Pretrained and Learnable Speaker Representations for Multi-Speaker Multi-Style Text-to-Speech,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2021, pp. 8588–8592.
- [10] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR Corpus Based on Public Domain Audio Books,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2015, pp. 5206–5210.
- [11] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks,” in *International Conference on Machine Learning*, 2006, pp. 369–376.
- [12] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust Speech Recognition via Large-Scale Weak Supervision,” in *International Conference on Machine Learning*, 2023, pp. 28 492–28 518.
- [13] J. Kim, J. Kong, and J. Son, “Conditional Variational Autoencoder with Adversarial Learning for End-to-End Text-to-Speech,” in *International Conference on Machine Learning*, 2021, pp. 5530–5540.
- [14] B. Biggio, B. Nelson, and P. Laskov, “Poisoning Attacks against Support Vector Machines,” in *International Conference on Machine Learning*, 2012, p. 1467–1474.
- [15] 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.
- [16] 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.
- [17] A. Nagrani, J. S. Chung, and A. Zisserman, “VoxCeleb: A Large-Scale Speaker Identification Dataset,” in *Interspeech 2017*, 2017, pp. 2616–2620.
- [18] T. Thebaud, S. Joshi, H. Li, M. Sustek, J. Villalba, S. Khudanpur, and N. Dehak, “Clustering Unsupervised Representations as Defense Against Poisoning Attacks on Speech Commands Classification System,” in *IEEE Automatic Speech Recognition and Understanding Workshop*, 2023, pp. 1–8.