

# NTU-NPU System for Voice Privacy 2024 Challenge

Nikita Kuzmin<sup>1,2</sup>, Hieu-Thi Luong<sup>1</sup>, Jixun Yao<sup>3</sup>, Lei Xie<sup>3</sup>, Kong Aik Lee<sup>4</sup>,  
Eng Siong Chng<sup>1</sup>

<sup>1</sup>Nanyang Technological University <sup>2</sup>Institute for Infocomm Research, A\*STAR

<sup>3</sup>Audio, Speech and Language Processing Group (ASLP@NPU)

<sup>4</sup>The Hong Kong Polytechnic University

s220028@e.ntu.edu.sg, hieuthi.luong@ntu.edu.sg, yaojx@mail.nwpu.edu.cn

## Abstract

In this work, we describe our submissions for the Voice Privacy Challenge 2024. Rather than proposing a novel speech anonymization system, we enhance the provided baselines to meet all required conditions and improve evaluated metrics. Specifically, we implement emotion embedding and experiment with WavLM and ECAPA2 speaker embedders for the B3 baseline. Additionally, we compare different speaker and prosody anonymization techniques. Furthermore, we introduce Mean Reversion F0 for B5, which helps to enhance privacy without a loss in utility. Finally, we explore disentanglement models, namely  $\beta$ -VAE and NaturalSpeech3 FACodec.

**Index Terms:** Voice Privacy Challenge 2024, Speaker Anonymization, Emotion Embedding, Voice Conversion

## 1. Introduction

Voice Privacy Challenge 2024 [1] introduces a new metric to evaluate the ability to maintain the conveyed emotion of the original source speaker.

Given the baselines provided by organizers, we test various techniques and methods aimed at improving the evaluated metrics. Specifically, we create submissions for all four of EER conditions by using a modified version of B3 [2] and B5 [3] systems, with the former responsible for conditions with EER lower than 30% and the latter for those with EER from 30% and above. Additionally, we experiment with disentanglement-based models such as  $\beta$ -VAE [4] and NaturalSpeech3 FACodec [5].

The rest of the paper is constructed as follows: Section 2 summarizes the two baselines that we used for our submissions, Section 3 introduces all the modifications and proposed techniques used in our experiments, Section 4 provides the detailed results and the systems that we submitted for evaluation, Section 5 concludes our findings. Furthermore, we provide the summary tables with a description of submitted systems in Appendix A.

## 2. Baseline Systems

In this section, we summarize two baseline systems, B3 and B5, that we use in our experiments. A modified version of B3 is used to create submission for the first condition (EER<sub>1</sub>) with EER between 10% and 20% and the second condition (EER<sub>2</sub>) with EER between 20% and 30%, while a modified version of B5 is used for the third condition (EER<sub>3</sub>) with EER between 30% and 40% and the last condition (EER<sub>4</sub>) with EER larger than 40%. The details about all proposed changes are laid out in Section 3.

### 2.1. B3

The baseline system B3 uses a Wasserstein generative adversarial network with Quadratic Transport Cost (WGAN-QC) [6] to generate artificial pseudo-speaker embeddings, anonymizing the speaker's identity through four main steps:

- 1. Phonetic Transcriptions Extraction:** Phonetic transcriptions are extracted using an end-to-end automatic speech recognition (ASR) model with a hybrid CTC-attention architecture.
- 2. Speaker Embeddings** Speaker embeddings are obtained using an adapted Global Style Tokens (GST) model [7].
- 3. Anonymization:** The original speaker embedding is swapped with an artificial one generated by a WGAN. If the cosine distance between the artificial and original embeddings is less than 0.3, the replacement is considered successful. Otherwise, the process is repeated up to 30 times. Additionally, the pitch and energy values for each phoneme are adjusted using random values between 0.6 and 1.4.
- 4. Speech Synthesis:** The anonymized speaker embedding, adjusted prosody, and original phonetic transcription are used to create anonymized speech using the FastSpeech2 model and HiFi-GAN [8] vocoder, as implemented in IMS-Toucan [9].

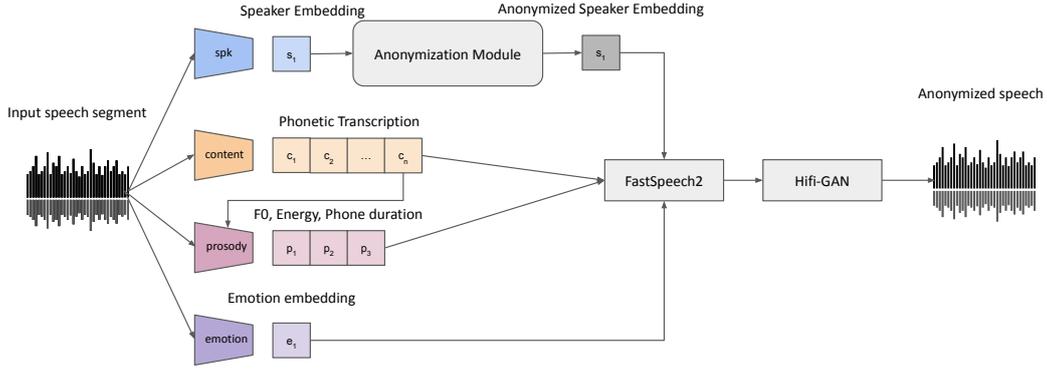


Figure 1: Schematic diagram of Modified system B3.

## 2.2. B5

The B5 system used a HiFi-GAN model conditioned on fundamental frequency and a linguistic representation of the source utterance along with speaker embedding of a designated speaker to generate anonymized speech.

1. **Fundamental Frequency (F0):** B5 uses a pytorch implementation of YAAPT Pitch Tracking [10] to extract F0 from speech. In authors' of B5 thesis [3], the original authors of B5 suggest several complementary normalization or transformations to be applied to F0, none of which are included in B5.
2. **Linguistic Representation:** B5 uses the output of a vector quantization bottleneck layer (VQ-BN) put the top of the acoustic model (AM) of an automatic speech recognition (ASR) trained to classify left-biphone as the linguistic representation.
3. **Speaker Embedding:** a designed speaker embedding of a speaker included in the training stage is used to change the voice of anonymized speech. We randomly pick a speaker embedding to anonymize each utterance similar to the B5 baseline provided by the organizer.

We use the same pre-trained B5 model provided by organizers without doing any further training or tuning. Instead, we introduce a new method of transformation that can be applied to F0 in the inference stage, which is discussed in Section 3.3

## 3. Our Methods

In this subsection, we elaborate on the details of proposed modifications to the baseline systems.

### 3.1. Modifications of B3

We experimented with the following main modifications for this system is as follows:

- Emotion embeddings are implemented as an additional input to the FastSpeech2 model.
- The Global Style Tokens (GST) model is replaced by different speaker embedders such as WavLM [11] and ECAPA2 [12].
- Some experiments related to speaker selection strategy for anonymization and prosody manipulation.

We start off with emotion embeddings. For extracting emotion-based embeddings, we employ a fine-tuned Wav2Vec2 Large Robust model [13] on MSP-Podcast [14]. Notably, the model is pruned from 24 to 12 transformers, and the CNN component is frozen prior to fine-tuning. Embeddings are extracted from the hidden layer, which has 1024-dimensional vectors as output. We employ it to FastSpeech2 in the same way as speaker embeddings in [15] by adding one more linear projection and concatenating it with an output from Conformer [16].

In addition to the GST model, we implement different speaker embedding models such as ECAPA2 and WavLM with 128 and 512 embedding sizes correspondingly. As both these speaker embedders work on audios instead of spectrograms, we add a pre-trained HiFi-GAN to the setup for the second phase of FastSpeech2 training in order to generate audio and extract embeddings for cycle consistency loss. In contrast to the GST model, ECAPA2 and WavLM speaker embedders are frozen during FastSpeech2 training. Similarly, the HiFi-GAN is also frozen.

Furthermore, we explore various anonymization strategies, such as random speaker selection, which involves replacing the source speaker embedding for each utterance with a randomly selected embedding from a pool of embeddings. Additionally, we evaluated the importance of the usage of cross-gender for anonymization for the modified model. In cross-gender anonymization technique, we select a target speaker from the pool that has the opposite gender

with respect to the source speaker. Finally, we examine how different powers of prosody anonymization affect privacy-based and utility-based metrics. To be more specific, we experiment with different offsets for pitch and energy multipliers.

The training process is the same as for the B3 baseline system. We retrain each part of the system except HiFi-GAN and ASR model.

### 3.2. Disentanglement-based models

Next, we explore disentanglement-based models, which might be useful for removing speaker-related information from other components (such as prosody, content, and acoustic information). We compare two models:  $\beta$ -VAE [4] and NaturalSpeech3 FACodec [5]. Anonymization for these models is implemented by passing anonymized speaker embedding instead of source embedding and then performing voice conversion.

### 3.3. Mean Reversion F0 for the B5 system

For conditions EER<sub>3</sub> and EER<sub>4</sub>, we base our submissions on the B5 [3] system provided by the organizer. Champion proposes a voice anonymization model using a HiFi-GAN vocoder which takes in the F0 and the linguistic representation of a source utterance and then turns it into the speech with the voice of a target speaker. In his thesis, the author explores different transformation techniques that can be applied to F0 including linear transformation, Additive White Gaussian Noise, and quantization. In this work, we propose a new type of transformation that uses the original and the n-frame moving average F0 ( $F_0$ ), with n=32 in our calculation:

$$\hat{F}_0 = (1 - \alpha)F_0 + \alpha\overline{F_0} \quad (1)$$

with  $\alpha = 0$ , we get the original F0, and with  $\alpha = 1$ , we get the moving average F0. For any  $\alpha$  between 0 and 1, we obtain a mean reversion F0 which is a value weighted toward the mean. The motivation for this method is that we can reduce the dynamic range of F0, which is one characteristic of voice, and move it to a more neutral value. Moreover, the calculation is based on a short window instead of the entire utterance like the linear transformation method. Note that, we remove unvoiced frames when calculating the moving average F0.

We then apply an Additive White Gaussian Noise on top of mean reversion F0 to push up the EER. Figure 2 shows an example of the mean reversion F0 with and without additive noise.

## 4. Experiments

In this section, we provide experimental results for the system modifications explained in Section 3. To begin, we present the results for different modifications of system B3. The modified system B3 is illustrated in Figure 1.

### 4.1. Modified B3

From the results in 1, it can be observed that emotion embeddings help to improve Emotion Recognition performance while maintaining ASR performance at the same level. However, there is some degradation in privacy, which might be due to speaker identity leakage in the emotion embeddings. In addition, we provide experimental results for this system without prosody anonymization to check how modifications of prosody affect SER performance. As shown in the results, removing prosody modifications improves SER and ASR but also reduces privacy, making this system suitable for condition with a minimum EER<sub>1</sub> = 10%.

Table 2 shows a comparison between WGAN and Random-Speaker anonymization techniques. There is almost no difference in the privacy and utility metrics for these methods, so we chose to stick with Random-Speaker as it requires no training.

Table 2: Comparison between WGAN anonymization strategy trained on LibriTTS-clean-100 and Random Speaker (Rnd-Spk) selection from LibriTTS-train-clean-100 per each source utterance.

Anon Type	EER		UAR		WER	
	dev	test	dev	test	dev	test
WGAN	25.20	27.78	38.40	37.70	4.30	4.40
Rnd-Spk	<b>25.76</b>	<b>28.42</b>	37.97	37.39	4.33	4.33

Next, Table 3 compares different ranges for multipliers involved in prosody anonymization. The results indicate that fewer prosody modifications result in worse privacy but better utility. This finding is useful for VPC2024 as it allows us to find a better trade-off between privacy and utility for specific EER conditions.

Table 3: Comparison between difference range for F0 and energy multipliers. The bottom row corresponds to the system without prosody manipulation.

Multiplier Range	EER		UAR		WER	
	dev	test	dev	test	dev	test
[0.6, 1.4]	25.76	28.42	37.97	37.39	4.33	4.33
[0.7, 1.3]	23.93	25.62	37.49	37.59	4.07	4.05
[0.8, 1.2]	22.70	25.92	38.01	37.96	3.89	3.91
[0.9, 1.1]	19.88	22.62	39.03	37.17	3.80	3.77
-	19.47	21.82	38.91	38.11	3.70	3.75

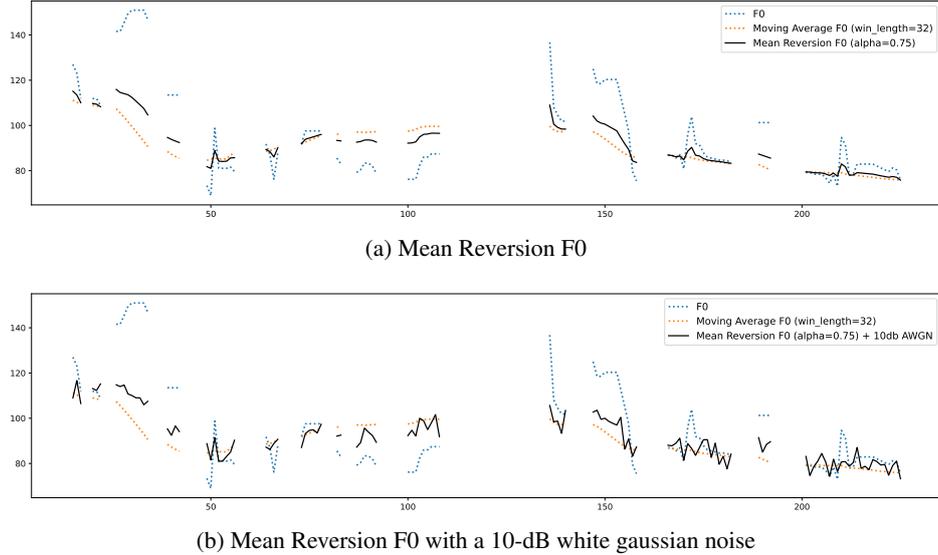


Figure 2: Examples of Mean Reversion F0 with and without additive noise

Table 1: Comparison between systems with and without emotion embedding with different speaker embedder and prosody anonymization. + in speaker anonymization column corresponds to Random-Speaker selection from LibriTTS-clean-100 for each source utterance. + in prosody anonymization column corresponds to the systems with prosody multipliers from [0.6, 1.4] range.

Speaker Anonymization	Speaker Embedder	Prosody Anonymization	Emotion Embedding	EER		UAR		WER	
				dev	test	dev	test	dev	test
-	-	-	-	5.72	4.59	69.08	71.06	1.80	1.85
+	GST	+	-	25.76	28.42	37.97	37.39	4.33	4.33
+	GST	+	+	22.59	24.09	42.52	41.74	4.39	4.40
+	GST	-	+	16.88	17.45	42.76	43.21	<b>3.81</b>	<b>3.83</b>
+	WavLM	-	+	17.97	16.64	<b>43.84</b>	<b>45.67</b>	4.54	4.69
+	ECAPA2	-	+	19.48	22.55	42.53	42.37	4.83	4.80

Table 4: Comparison between  $\beta$ -VAE and NS3 disentanglement models.

Model	EER		UAR		WER	
	dev	test	dev	test	dev	test
Original	5.72	4.59	69.08	71.06	1.80	1.85
$\beta$ -VAE	10.71	10.49	30.38	31.28	67.72	65.5
NS3	9.29	8.78	51.64	52.89	2.97	2.77

#### 4.2. Disentanglement-based models

The comparison results between  $\beta$ -VAE and NaturalSpeech3 FACodec are shown in Table 4. It can be seen that  $\beta$ -VAE performs poorly in utility-based tasks, likely because of the fact that content representations are not rich enough.

As one might notice from the results in Table 4, NaturalSpeech3 has decent utility results. Therefore,

we decided to employ anonymization techniques to improve privacy protection for NS3, aiming to meet a condition with minimum  $EER_1 = 10\%$ . We experimented with the following tricks: Additive White Gaussian Noise (AWGN) to Speaker Embedding and conversing a source speaker to a target speaker of the opposite gender (cross-gender). The results are shown in Table 5. As we can see, cross-gender conversion helps to improve privacy and ASR performance on the corresponding test sets. Interestingly, it also improves SER performance on both development and test sets. As expected, AWGN enhances privacy at the cost of utility.

Our results underscore the need for a balance between privacy and utility, as methods like AWGN and prosody anonymization can strengthen privacy but also impact system performance. This balance is essential for creating anonymization techniques that are both se-

Table 5: Comparison between NaturalSpeech3 FA-Codec systems with different power of AWGN applied to speaker embedding. The speaker anonymization module consists of averaging 100 embeddings randomly selected from a pool of 200 farthest embeddings (LibriTTS-train-clean-100) from source utterance by cosine scoring.

scale, $10e^{-3}$	Cross Gender	EER		UAR		WER	
		dev	test	dev	test	dev	test
-	-	9.29	8.78	51.64	52.89	2.97	2.77
75	-	12.25	9.14	48.00	48.09	4.66	4.63
75	+	12.09	10.46	49.20	49.12	4.97	4.60
78	+	12.42	10.24	49.10	49.39	5.44	5.07
80	-	12.63	9.42	47.33	48.35	5.37	5.40
80	+	13.66	10.10	48.82	48.95	5.69	5.37
90	-	12.41	10.45	47.61	47.10	7.04	6.45

Table 6: Evaluation results of B5 using Mean Reversion F0 with different values of  $\alpha$  in inference stage

$\alpha$	EER		UAR		WER	
	dev	test	dev	test	dev	test
0.00	31.64	31.36	39.18	38.24	4.79	4.44
0.25	32.13	32.03	<b>39.61</b>	38.38	4.74	4.54
0.50	33.48	34.08	38.60	37.34	<b>4.62</b>	4.54
<b>0.75</b>	<b>38.56</b>	37.48	38.06	37.60	4.70	4.47
1.00	37.91	<b>37.93</b>	38.50	<b>38.78</b>	4.79	<b>4.43</b>

cure and functional.

### 4.3. Modified B5

Table 7: Evaluation results of B5 using Mean Reversion F0 ( $\alpha = 0.75$ ) and AWGN with different magnitude of noise in inference stage

dB	EER		UAR		WER	
	dev	test	dev	test	dev	test
0	38.56	37.48	38.06	37.60	4.70	4.47
5	39.58	40.00	38.91	37.12	4.67	4.49
<b>10</b>	42.46	<b>43.15</b>	<b>39.41</b>	<b>38.47</b>	<b>4.63</b>	<b>4.40</b>
15	<b>42.97</b>	40.36	38.50	37.49	4.66	4.50
30	41.43	39.62	38.41	37.88	4.77	4.64

Table 6 lists the results of the Mean Reversion F0 method discussed in Section 3.3 given different  $\alpha$  values. We can see a general trend that EER increases when  $\alpha$  is increased while UAR and WER fluctuate but not very significantly. We submitted the sample generated with  $\alpha = 0.75$  for the condition EER<sub>3</sub>.

For the last condition EER<sub>4</sub>, we add a 10-dB AWGN to the mean reversion F0 with  $\alpha = 0.75$  and manage to obtain an EER above 40%. The result can be found in Table 8.

We note that the EER results of these systems were highly volatile during our experiments, often produc-

ing different results even if we ran with the same configuration. It seems that convergence of an attacker ASV model depends on factors such as the machine, GPU, randomly picked speaker embedding, and other random parameters. The systems that we selected for submissions were based on the results available at that time and represented our methods of Mean Reversion F0 and AWGN.

### 4.4. Submitted systems

In this subsection, we provide a summary of all submitted systems. Table 8 shows privacy and utility results for each of the conditions.

Table 8: Results summary for all submitted systems grouped by achieved privacy conditions.

Condition	System ID	EER		UAR		WER	
		dev	test	dev	test	dev	test
EER <sub>1</sub>	1a	12.09	10.46	49.20	49.12	4.97	4.60
	1b	16.88	17.45	42.76	43.21	3.81	3.83
EER <sub>2</sub>	2a	21.47	24.13	44.67	42.78	4.21	4.29
	2b	20.07	22.85	39.18	37.67	3.61	3.68
EER <sub>3</sub>	3	38.56	37.48	38.06	37.60	4.70	4.47
EER <sub>4</sub>	4	42.46	43.15	39.41	38.47	4.63	4.40

Additionally, we prepared the tables in the Appendix A with a summary of architecture, input, output values, and training data for components in submitted systems.

## 5. Conclusion

In this system description, we modified the baseline systems (B3 and B5) for the Voice Privacy Challenge 2024 to enhance speaker anonymization while preserving emotional and content features. Specifically, we integrated emotion embeddings and different speaker embedders such as WavLM and ECAPA2 into system B3. In addition, we explored random-speaker and cross-gender anonymizations and different setups of prosody manipulation. For B5, we introduced Mean Reversion method and AWGN for prosody which allowed us to enhance privacy while maintaining utility. Finally, we experimented with disentanglement-based approaches such as  $\beta$ -VAE and NaturalSpeech3. An additional analysis on NaturalSpeech3 showed promising results for the Voice Privacy problem.

## 6. Acknowledgements

The computational work for this article was partially performed on resources of the National Supercomputing Centre, Singapore.

## 7. References

- [1] 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," 2024.
- [2] S. Meyer, F. Lux, J. Koch, P. Denisov, P. Tilli, and T. Vu, "Prosody is not identity: A speaker anonymization approach using prosody cloning," 06 2023, pp. 1–5.
- [3] P. Champion, "Anonymizing speech: Evaluating and designing speaker anonymization techniques," Ph.D. dissertation, Université de Lorraine, 2023.
- [4] H. Lu, D. Wang, X. Wu, Z. Wu, X. Liu, and H. Meng, "Disentangled speech representation learning for one-shot cross-lingual voice conversion using  $\beta$ -vae," in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 814–821.
- [5] Z. Ju, Y. Wang, K. Shen, X. Tan, D. Xin, D. Yang, Y. Liu, Y. Leng, K. Song, S. Tang, Z. Wu, T. Qin, X.-Y. Li, W. Ye, S. Zhang, J. Bian, L. He, J. Li, and S. Zhao, "Naturalspeech 3: Zero-shot speech synthesis with factorized codec and diffusion models," *ArXiv*, vol. abs/2403.03100, 2024.
- [6] H. Liu, X. Gu, and D. Samaras, "Wasserstein gan with quadratic transport cost," *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4831–4840, 2019.
- [7] Y. Wang, D. Stanton, Y. Zhang, R.-S. Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," in *International conference on machine learning*. PMLR, 2018, pp. 5180–5189.
- [8] 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.
- [9] F. Lux, J. Koch, S. Meyer, T. Bott, N. Schauffler, P. Denisov, A. Schweitzer, and N. T. Vu, "The IMS Toucan System for the Blizzard Challenge 2023," in *Blizzard Challenge Workshop*. ISCA Speech Synthesis SIG, 2023.
- [10] S. A. Zahorian and H. Hu, "A spectral/temporal method for robust fundamental frequency tracking," *The Journal of the Acoustical Society of America*, vol. 123, no. 6, pp. 4559–4571, 2008.
- [11] 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, and F. Wei, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," 2021.
- [12] J. Thienpondt and K. Demuyne, "Speaker embeddings with weakly supervised voice activity detection for efficient speaker diarization," 05 2024.
- [13] J. Wagner, A. Triantafyllopoulos, H. Wierstorf, M. Schmitt, F. Burkhardt, F. Eyben, and B. W. Schuller, "Dawn of the transformer era in speech emotion recognition: Closing the valence gap," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 45, no. 09, pp. 10 745–10 759, sep 2023.
- [14] R. Lotfian and C. Busso, "Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings," *IEEE Transactions on Affective Computing*, vol. 10, pp. 471–483, 2019.
- [15] S. Meyer, P. Tilli, F. Lux, P. Denisov, J. Koch, and N. T. Vu, "Cascade of phonetic speech recognition, speaker embeddings gan and multispeaker speech synthesis for the VoicePrivacy 2022 Challenge," in *Proc. 2nd Symposium on Security and Privacy in Speech Communication*, 2022.
- [16] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," 10 2020, pp. 5036–5040.
- [17] R. Kumar, P. Seetharaman, A. Luebs, I. Kumar, and K. Kumar, "High-fidelity audio compression with improved rvqgan," *Advances in Neural Information Processing Systems*, vol. 36, 2024.

## A. Summaries for submitted systems

Table 9: *Description of Systems 1a.*

#	Module	Description	Output features	Data
1	Encoder [17]	4 Downsampling Convolution-based Layers with Snake activation function Input: speech waveform	Output vector <sup>256</sup>	Librilight train
2	Prosody extractor	Factorized Vector Quantization with 1 quantizer, codebook size: 1024	Prosody vector <sup>256</sup>	Librilight train
3	Content extractor	Factorized Vector Quantization with 2 quantizers, codebook size: 1024	Content vector <sup>256</sup>	Librilight train
4	Speaker embedding extractor	Several Conformer blocks	Speaker embedding <sup>256</sup>	Librilight train
5	Speaker anonymization module	Averaged 100 embeddings randomly selected from a pool of 200 farthest embeddings from source by cosine scoring AWGN with scale= 0.075 Cross-gender	Anonymized speaker embedding <sup>256</sup>	LibriTTS: train-clean-100
6	Decoder [17]	Upsampling Convolution-based Layers with Snake activation function	speech waveform	Librilight train

Table 10: Description of Systems 1b, 2a, 2b.

#	Module	Description	Output features	Data
1	Prosody extractor	Phone aligner: 6-layer CNN + LSTM with CTC loss F0 estimation using Praat F0, energy, durations normalized by each vector’s mean	F0 <sup>1</sup> , energy <sup>1</sup> phone durations <sup>1</sup>	LibriTTS: train-clean-100
2	ASR	End-to-end with hybrid CTC-attention Input: log mel Fbank <sup>80</sup> Encoder: Branchformer Decoder: Transformer CTC and attention criteria	phonetic transcript with pauses and punctuation	LibriTTS: train-clean-100 train-other-500
3	Speaker embedding extractor	GST, trained jointly with SS model Input: mel spectrogram <sup>80</sup> 6 hidden layers + 4-head attention	GST speaker embedding <sup>128</sup>	LibriTTS: train-clean-100
4	Emotion embedding extractor	<b>1b, 2a:</b> Dimensional Speech Emotion Recognition Model based on Wav2vec 2.0 Input: Wav2vec 2.0 Large features	emotion embedding <sup>1024</sup>	MSP-Podcast (v1.7)
		<b>2b:</b> –	–	–
5	Prosody modification module	<b>1b, 2b:</b> –	–	–
		<b>2a:</b> Value-wise multiplication of F0 and energy with random values in [0.7, 1.3)	F0 <sup>1</sup> , energy <sup>1</sup>	LibriTTS: train-clean-100
6	Speaker anonymization module	<b>1b:</b> Averaged 100 embeddings randomly selected from a pool of 200 farthest embeddings from source by cosine scoring + cross-gender <b>2a, 2b:</b> Random Speaker selection per each source utterance + cross-gender	Anonymized speaker embedding <sup>128</sup>	LibriTTS: train-clean-100
7	SS model	IMS Toucan implementation of FastSpeech2 Input: F0 <sup>1</sup> + energy <sup>1</sup> + phone durations <sup>1</sup> + phonetic transcript + GST embeddings <sup>128</sup> ( <b>1b, 2a:</b> + emotion embeddings <sup>1024</sup> ) Training criterion defined in FastSpeech2	mel spectrogram <sup>80</sup>	LibriTTS: train-clean-100
8	Vocoder	HiFi-GAN vocoder Input: mel spectrogram <sup>80</sup> Training criterion defined in HiFi-GAN	speech waveform	LibriTTS: train-clean-100

Table 11: Description of Systems 3 and 4.

#	Module	Description	Output features	Data
1	F0 extractor	F0 extracted with s pytorch implementation of YAAPT <b>3:</b> Using Mean Reversion F0 ( $\alpha = 0.75$ ) in inference <b>4:</b> Using Mean Reversion F0 ( $\alpha = 0.75$ ) and 10-db AWGN	F0	N/A
2	ASR AM with VQ	Acoustic Model trained to identify left bi-phones and a VQ bottleneck layer	Linguistic representation	VoxPopuli Librispeech: train-clean-100
3	Speaker embedding	One-hot vector represented speaker in training set	Speaker embedding	LibriTTS: train-clean-100
4	Speech Synthesis	HiFi-GAN vocoder Input: F0 + linguistic representation + speaker embedding	Speech waveform	LibriTTS: train-clean-100