

QR-VC: Leveraging Quantization Residuals for Linear Disentanglement in Zero-Shot Voice Conversion

Youngjun Sim*

*Graduate School of Artificial
Intelligence
POSTECH*

Pohang, South Korea

youngjunsim@postech.ac.kr

Jinsung Yoon*

*Graduate School of Artificial
Intelligence
POSTECH*

Pohang, South Korea

truestar2001@postech.ac.kr

Wooyeol Jeong

*Graduate School of Artificial
Intelligence
POSTECH*

Pohang, South Korea

jungwy0106@postech.ac.kr

Young-Joo Suh†

*Graduate School of Artificial
Intelligence
POSTECH*

Pohang, South Korea

yjsuh@postech.ac.kr

Abstract—Zero-shot voice conversion is a technique that alters the speaker identity of an input speech to match a target speaker using only a single reference utterance, without requiring additional training. Recent approaches extensively utilize self-supervised learning features with K-means quantization to extract high-quality content representations while removing speaker identity. However, this quantization process also eliminates fine-grained phonetic and prosodic variations, degrading intelligibility and prosody preservation. While prior works have primarily focused on quantized representations, quantization residuals remain underutilized and deserve further exploration. In this paper, we introduce a novel approach that fully utilizes quantization residuals by leveraging temporal properties of speech components. This facilitates the disentanglement of speaker identity and the recovery of phonetic and prosodic details lost during quantization. By applying only K-means quantization and linear projections, our method achieves simple yet effective disentanglement, without requiring complex architectures or explicit supervision. This allows for high-fidelity voice conversion trained solely with reconstruction losses. Experiments show that the proposed model outperforms existing methods across both subjective and objective metrics. It achieves superior intelligibility and speaker similarity, along with improved prosody preservation, highlighting the impact of our Linear Disentangler module.

Index Terms—voice conversion, self-supervised learning, quantization, linear disentanglement

I. INTRODUCTION

Zero-shot voice conversion (VC) aims to transform the speaker identity of a source speech into that of an arbitrary target using only a single utterance, without additional training. Generating high-fidelity converted speech requires a distinct disentanglement of content-related and speaker-specific components while preserving intelligibility, prosody, and speaker similarity.

Self-supervised learning (SSL) models [1–3] have gained significant attention in VC research as they are trained to encode latent speech representations. The resulting SSL features have been shown to linearly predict various speech

attributes [3–5], suggesting that key components of speech are linearly separable. This separability underpins the design of our linear disentanglement module.

Over the past few years, zero-shot VC performance has been greatly enhanced by integrating SSL features with techniques such as diffusion models [6], conformer-based architectures [7], normalizing flows [8], explicit supervision like prosodic labels [7, 9], and the use of external speaker encoders [8, 9], within a disentanglement module. Nevertheless, these approaches typically rely on complex deep architectures, resulting in increased model size and computational overhead.

One key property of SSL features is that nearby features share phonetic information [10, 11]. Leveraging this property, K-means quantization (KQ) has been widely adopted as a strong content encoder in various VC models [7, 9, 12]. This process effectively removes speaker information while capturing essential content representations.

However, quantization eliminates fine-grained phonetic and prosodic variations, referred to as speaking variations, leading to intelligibility degradation and poor prosody preservation. Existing models struggle to reconstruct these lost details, suggesting that quantization residuals are worth exploring further.

We propose QR-VC, a novel zero-shot VC model that leverages quantization residuals and the temporal properties of speech components to construct a simple yet effective architecture. Our model disentangles phonetic content, speaker identity, and speaking variations using only K-means quantization and linear projection layers. This approach enables high-fidelity VC relying solely on reconstruction losses, without requiring explicit supervision or complex modules. In particular, the proposed Linear Disentangler module, composed of time-invariant and time-variant linear bottlenecks, addresses the limitations of quantization by restoring fine-grained speaking variations, even with a small codebook. Through extensive experiments, we demonstrate that QR-VC outperforms existing methods in terms of intelligibility, speaker similarity, prosody preservation, and naturalness. Our

*Equal contribution

†Corresponding author

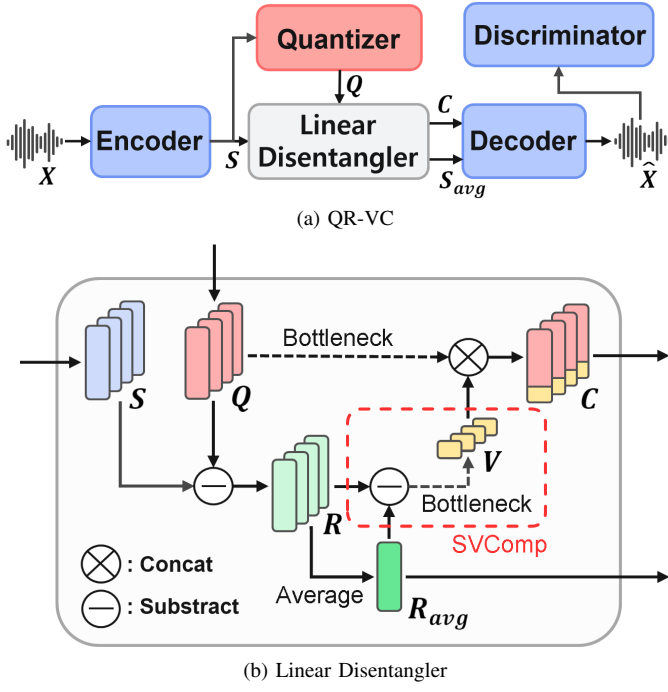


Fig. 1. Overview of QR-VC. (a) Overall architecture. (b) Detailed structure of the Linear Disentangler module.

results highlight the importance of quantization residuals and validate the effectiveness of a linear disentanglement framework in zero-shot VC. Audio samples are available at <https://simyoungjun.github.io/QR-VC-DEMO/>

II. METHOD

A. Overview

QR-VC consists of a WavLM [3]-based encoder, a Linear Disentangler, and a HiFi-GAN [14] decoder. The key improvement lies in leveraging quantization residuals to disentangle speaker identity and speaking variations through a linear structure. This approach preserves fine-grained prosodic and phonetic details while ensuring speaker similarity.

B. Encoder and Quantizer

The encoder extracts SSL features S from the 6th layer of WavLM given a speech input X , defined as:

$$S = [s_1, s_2, \dots, s_T] \in \mathbb{R}^{1024 \times T}. \quad (1)$$

where T denotes the number of frames. Previous studies have demonstrated that these features, in particular, retain rich phonetic, prosodic, and speaker information [5, 10].

To obtain a discrete content representation Q , we apply K-means quantization using a codebook E :

$$\mathbf{E} = \{e_1, e_2, \dots, e_K\} \in \mathbb{R}^{K \times 1024}. \quad (2)$$

where K is the number of codebook entries. Each frame s_t is mapped to the nearest entry:

$$q_t = \arg \min_{e_k \in E} \|s_t - e_k\|_2^2, \quad Q = [q_1, q_2, \dots, q_T]. \quad (3)$$

Quantization removes speaker identity while preserving essential linguistic content. However, fine-grained speaking variations, such as phonetic and prosodic details, are also lost.

C. Linear Disentangler

To recover these lost details, we introduce a disentanglement mechanism that decomposes the quantization residual R :

$$r_t = s_t - q_t, \quad R \in \mathbb{R}^{1024 \times T}. \quad (4)$$

This quantization residual R preserves fine-grained phonetic and prosodic details, along with speaker information, as shown in Figure 2. The disentangler further separates these components into time-invariant speaker identity and time-varying speaking variations.

1) *Speaker Identity Extraction*: Since speaker identity is generally consistent across an utterance, we extract it by computing the temporal average of R :

$$R_{avg} = \frac{1}{T} \sum_{t=1}^T r_t. \quad (5)$$

This removes time-varying fluctuations while preserving time-invariant speaker identity characteristics.

2) *Speaking Variation Compensation (SVComp)*: To isolate speaking variations V , we subtract the speaker identity vector R_{avg} from each quantization residual frame r_t . The result is passed through a 1×1 convolutional bottleneck layer, which serves as a linear projection. This operation captures important time-varying information:

$$V = \text{Bottleneck}(R - R_{avg}), \quad v_t \in \mathbb{R}^M. \quad (6)$$

where M is the bottleneck dimension. The resulting representation V retains prosodic and phonetic details while suppressing speaker information.

3) *Final Content Representation*: The disentangled content representation C is obtained by concatenating the bottlenecked quantized representation Q' with the extracted speaking variations V . The two components have dimensions of $(1024 - M) \times T$ and $M \times T$, respectively, forming a 1024-dimensional representation:

$$C = Q' \otimes V, \quad c_t \in \mathbb{R}^{1024}. \quad (7)$$

This process compensates for the loss of speaking variations in quantization, leading to improved content and prosody preservation.

D. Decoder and Training Objective

The decoder G reconstructs the speech waveform \hat{X} using HiFi-GAN [14] generator, taking the concatenated content C and speaker identity R_{avg} as the input:

$$\hat{X} = G(C + R_{avg}). \quad (8)$$

TABLE I
SUBJECTIVE AND OBJECTIVE EVALUATION RESULTS FOR MOS, SMOS, WER, CER, EER, SECS, E-PCC, AND F0-PCC.

Model	MOS↑	SMOS↑	WER(%)↓	CER(%)↓	EER(%)↓	SECS↑	E-PCC↑	F0-PCC↑
SSR-VC [9]	2.97±0.17	3.04±0.16	11.80	4.64	18.6	0.725	0.701	0.692
SEF-VC [7]	3.98±0.14	3.41±0.15	7.57	2.47	16.6	0.707	0.708	0.657
Phoneme Hallucinator [13]	3.81±0.18	3.76±0.13	10.78	4.76	12.8	0.743	0.793	0.716
DDDM-VC [6]	3.69±0.20	4.03±0.18	12.05	5.19	8.0	0.791	0.685	0.685
kNN-VC [10]	2.78±0.15	3.68±0.19	37.58	20.81	4.0	0.798	0.622	0.628
QR-VC (Ours)	4.14±0.12	4.27±0.13	5.57	2.07	5.0	0.815	0.885	0.719

The model is trained with a combination of adversarial losses $L_{adv}(G)$ and $L_{adv}(D)$ for the decoder G and the discriminator D , a feature matching loss L_{fm} , and mel-spectrogram reconstruction loss L_{mel} from HiFi-GAN:

$$\begin{aligned} L_G &= L_{adv}(G) + \lambda_{fm} L_{fm} + \lambda_{mel} L_{mel}, \\ L_D &= L_{adv}(D). \end{aligned} \quad (9)$$

Additionally, to improve computational efficiency and enhance model generalization, we randomly extract segments of latent representations and corresponding audio segments from ground truth raw waveforms as training targets.

E. Inference

During inference, the model converts speech by replacing the speaker representation. Given a source utterance and a single target utterance, the disentangler extracts content C^{src} from the source and speaker identity S_{avg}^{tgt} from the target. The decoder then generates the converted waveform \hat{X}_{conv} :

$$\hat{X}_{conv} = G(C^{src} + R_{avg}^{tgt}). \quad (10)$$

III. EXPERIMENTS

A. Experimental Setups

We use the LibriSpeech dataset [15], specifically the train-clean-100 subset for training, which consists of 100 hours of speech, and the test-clean subset for evaluation, which includes 5.4 hours of speech from 40 speakers. All audio is downsampled to 16 kHz. Preprocessing settings include FFT, window, and hop sizes of 1280, 1280, and 320, respectively. KQ is applied using mini batch k-means [16] with a batch size of 2048 and $K = 256$ clusters. The speaking variations V is reduced to $M = 8$ dimensions, and the bottlenecked quantized representation Q' has a dimension of 1016. The length of the randomly extracted segments, as defined in section II-D, is 28. The encoder is frozen during training, while other modules are optimized with AdamW with $\beta_1 = 0.8$, $\beta_2 = 0.99$, weight decay $\lambda = 0.01$, and an initial learning rate of 2×10^{-4} , decaying by a factor of 0.999 each epoch, with a batch size of 32.

B. Baseline Methods

The proposed method is compared against several state-of-the-art zero-shot VC approaches. SSR-VC [9] employs HuBERT with a 100-codebook KQ as the encoder with external

speaker encoder and incorporates a pitch predictor with a dedicated loss function. SEF-VC [7] integrates HuBERT with a 2000-codebook KQ and a conformer-based network, incorporating a supervised prosody predictor to enhance prosody modeling. kNN-VC [10] utilizes WavLM 6th layer features, replacing each source feature with the average of its nearest neighbors in a target utterance using k-nearest neighbors (kNN) regression. Phoneme Hallucinator [13] extends kNN-VC by incorporating a VAE-based architecture to generate synthetic target-matching features, improving phonetic coverage in zero-shot scenarios. DDDM-VC [6] employs WavLM features alongside a diffusion-based network for disentangling speaker identity and content information.

C. Evaluation Metrics

We conduct both subjective and objective evaluations. In the subjective evaluation, 50 participants rate naturalness and similarity on a 5-point scale using the mean opinion score (MOS) and similarity mean opinion score (SMOS), with 95% confidence intervals. Experiments are conducted on 20 unseen speakers from the LibriTTS test set. For the objective evaluation, we used 6 metrics. Intelligibility was assessed via word error rate (WER) and character error rate (CER) calculated with a Whisper-based ASR model¹ [17]. Speaker similarity was measured with equal error rate (EER) and speaker embedding cosine similarity (SECS) using a pre-trained speaker encoder, Resemblyzer². Prosody preservation was evaluated with the Pearson correlation coefficient (PCC) [18] for both F0 (F0-PCC) and energy (E-PCC), where higher PCC values indicate better prosody preservation. For evaluation, 1000 utterances were randomly selected from the Librispeech test set, and 500 unseen-to-unseen converted utterances are generated.

IV. RESULTS

A. Comparison with Baselines

Table I presents the subjective and objective evaluation results comparing the proposed method with baseline models. The proposed model achieves the lowest WER of 5.57% and CER of 2.07%, demonstrating superior intelligibility. It also obtains the highest speaker similarity, with a SMOS of 4.27 and a SECS of 0.815, and shows strong prosody preservation

¹<https://github.com/openai/whisper>

²<https://github.com/resemble-ai/Resemblyzer>

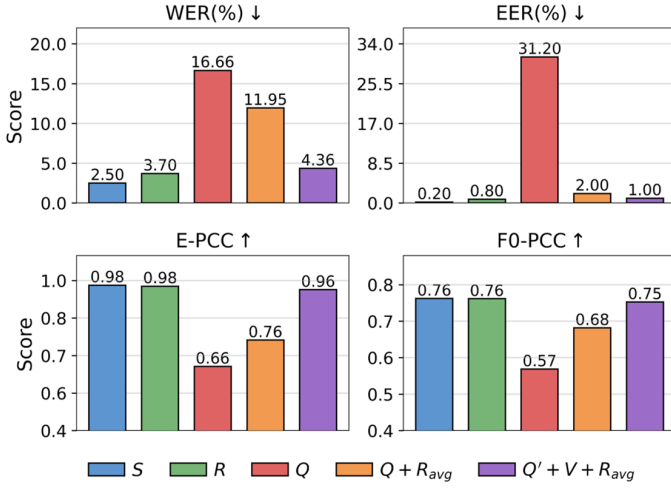


Fig. 2. Objective evaluation of reconstructed speech from different feature representations.

with an F0-PCC of 0.719 and an E-PCC of 0.885. Additionally, the MOS results indicate the best naturalness. These results confirm that QR-VC enables high-fidelity zero-shot voice conversion, preserving both content and prosody while enhancing speaker similarity. This validates that the Linear Disentangler ensures distinct separation of speech components, preventing interference or information leakage.

B. Feature Representation Analysis

We conduct a speech reconstruction task using representations derived from SSL features to evaluate their ability to recover the original waveform. This experiment assesses how well various speech components are preserved in the reconstructed speech, explicitly analyzing the information retained by each representation. Figure 2 presents the objective evaluation of reconstructed speech using various feature representations, including SSL features S , quantized representation Q , quantization residual R , bottlenecked quantized representation Q' , speaker identity R_{avg} , and speaking variations V .

SSL features S encode rich phonetic, prosodic, and speaker information, resulting in high intelligibility and speaker similarity. In contrast, quantized representation Q effectively removes speaker identity while preserving essential linguistic content, but also eliminates speaking variations, degrading prosody and content preservation. The quantization residual R retains fine-grained phonetic and prosodic details, along with speaker identity, suggesting its potential as a valuable feature for speech processing.

The combination of $Q + R_{avg}$ significantly improves the EER, confirming that R_{avg} successfully captures speaker identity. Moreover, it enhances reconstructed speech quality, reducing the WER. The combination of $Q' + V + R_{avg}$ achieves the best trade-off among intelligibility, prosody preservation, and speaker similarity. These results also validate the effectiveness of the Linear Disentangler’s ability to separate key speech components while preserving each of them.




Fig. 3. Comparison of QR-VC models with and without speaking variations compensation across different codebook sizes.

C. Effect of Speaking Variations Compensation (SVComp)

Fig. 3 shows the effect of SVComp across different codebook sizes. As the codebook size decreases, the performance gap between models with and without SVComp increases, suggesting that preserving fine-grained phonetic and prosodic details becomes increasingly difficult with lower quantization resolutions. Despite this, the model with SVComp consistently achieves lower WER while maintaining higher prosody consistency, as evidenced by the E-PCC and F0-PCC.

D. Ablation Studies

An ablation study is conducted to evaluate the contribution of each component in the proposed method, as shown in Table II. Removing SVComp (*w/o SVComp*) results in a substantial increase in WER while reducing F0-PCC and E-PCC, confirming its importance in preserving intelligibility and prosody. Eliminating the bottleneck layer in SVComp (*w/o bottleneck*) slightly improves intelligibility and prosody preservation but degrades speaker similarity, demonstrating its role in suppressing speaker information. Replacing the proposed speaker representation R_{avg} with an external speaker embedding [19] (*external spk emb*), in the absence of SVComp due to distribution mismatch, leads to a higher EER,

TABLE II
OBJECTIVE EVALUATION RESULTS OF ABLATION SYSTEMS.

Model	WER(%)↓	EER(%)↓	E-PCC↑	F0-PCC↑
Ours	5.57	5.0	0.885	0.719
-w/o SVComp	12.93	4.4	0.676	0.668
-w/o bottleneck	4.09	8.8	0.903	0.729
-external spk emb	8.55	22.6	0.836	0.664

indicating that the proposed method more effectively captures speaker identity.

These results validate the importance of each component in the proposed framework, demonstrating their critical role in preserving intelligibility, prosody, and speaker similarity.

V. CONCLUSION

This paper proposes QR-VC, a novel zero-shot voice conversion method that leverages quantization residuals in a linear disentanglement framework. By exploiting the temporal properties of speech components, the model disentangles linguistic content, speaker identity, and speaking variations using only K-means quantization and linear projections. This design enables high-fidelity voice conversion relying exclusively on reconstruction losses. Experiments show that QR-VC outperforms baselines in terms of naturalness, intelligibility, speaker similarity, and prosody preservation. Ablation studies confirm the importance of residual-based speaker representation and the bottleneck structure in disentangling speech components. Evaluations across codebook sizes highlight the effectiveness of the speaking variation compensation module, especially under low-resolution quantization. Future work includes utilizing the model’s simplicity for real-time and on-device voice conversion.

VI. ACKNOWLEDGMENT

This research was partly supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(2022R1A6A1A03052954), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.RS-2019-II191906, Artificial Intelligence Graduate School Program(POSTECH)), and the Korea Institute of Energy Technology Evaluation and Planning(KETEP) grant funded by the Ministry of Trade, Industry & Energy(MOTIE) of the Republic of Korea (No.20214810100010).

REFERENCES

- [1] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 29, pp. 3451–3460, 2021.
- [2] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [3] 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.
- [4] A. Pasad, J.-C. Chou, and K. Livescu, “Layer-wise analysis of a self-supervised speech representation model,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 914–921.
- [5] G.-T. Lin, C.-L. Feng, W.-P. Huang, Y. Tseng, T.-H. Lin, C.-A. Li, H.-y. Lee, and N. G. Ward, “On the utility of self-supervised models for prosody-related tasks,” in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 1104–1111.
- [6] H.-Y. Choi, S.-H. Lee, and S.-W. Lee, “Dddm-vc: Decoupled denoising diffusion models with disentangled representation and prior mixup for verified robust voice conversion,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 17 862–17 870.
- [7] J. Li, Y. Guo, X. Chen, and K. Yu, “Sef-vc: Speaker embedding free zero-shot voice conversion with cross attention,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 12 296–12 300.
- [8] J. Li, W. Tu, and L. Xiao, “Freevc: Towards high-quality text-free one-shot voice conversion,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [9] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhota, W.-N. Hsu, A. Mohamed, and E. Dupoux, “Speech resynthesis from discrete disentangled self-supervised representations,” *arXiv preprint arXiv:2104.00355*, 2021.
- [10] M. Baas, B. van Niekkerk, and H. Kamper, “Voice conversion with just nearest neighbors,” *arXiv preprint arXiv:2305.18975*, 2023.
- [11] E. Dunbar, N. Hamilakis, and E. Dupoux, “Self-supervised language learning from raw audio: Lessons from the zero resource speech challenge,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1211–1226, 2022.
- [12] X. Zhu, Y. Lv, Y. Lei, T. Li, W. He, H. Zhou, H. Lu, and L. Xie, “Vec-tok speech: Speech vectorization and tokenization for neural speech generation,” *arXiv preprint arXiv:2310.07246*, 2023.
- [13] S. Shan, Y. Li, A. Banerjee, and J. B. Oliva, “Phoneme hallucinator: One-shot voice conversion via set expansion,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 13, 2024, pp. 14 910–14 918.
- [14] 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.
- [15] 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.
- [16] O. Kramer and O. Kramer, “Scikit-learn,” *Machine learning for evolution strategies*, pp. 45–53, 2016.
- [17] 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*. PMLR, 2023, pp. 28 492–28 518.
- [18] I. Cohen, Y. Huang, J. Chen, J. Benesty, J. Benesty, J. Chen, Y. Huang, and I. Cohen, “Pearson correlation coefficient,” *Noise reduction in speech processing*, pp. 1–4, 2009.
- [19] S. Liu, Y. Cao, D. Wang, X. Wu, X. Liu, and H. Meng, “Any-to-many voice conversion with location-relative sequence-to-sequence modeling,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1717–1728, 2021.