



# Voice Conversion With Just Nearest Neighbors

Matthew Baas<sup>†</sup>, Benjamin van Niekerk<sup>†</sup>, Herman Kamper

MediaLab, Department of Electrical & Electronic Engineering, Stellenbosch University, South Africa  
20786379@sun.ac.za, benjamin.l.van.niekerk@gmail.com, kamperh@sun.ac.za

## Abstract

Any-to-any voice conversion aims to transform source speech into a target voice with just a few examples of the target speaker as a reference. Recent methods produce convincing conversions, but at the cost of increased complexity – making results difficult to reproduce and build on. Instead, we keep it simple. We propose k-nearest neighbors voice conversion (kNN-VC): a straightforward yet effective method for any-to-any conversion. First, we extract self-supervised representations of the source and reference speech. To convert to the target speaker, we replace each frame of the source representation with its nearest neighbor in the reference. Finally, a pretrained vocoder synthesizes audio from the converted representation. Objective and subjective evaluations show that kNN-VC improves speaker similarity with similar intelligibility scores to existing methods. Code, samples, trained models: <https://bshall.github.io/knn-vc>.

**Index Terms:** speech synthesis, voice conversion, k-nearest neighbors.

## 1. Introduction

The goal of voice conversion is to transform source speech into a target voice, keeping the content unchanged [1]. We can categorize voice conversion systems by how restrictive the set of source and target speakers are. The most general case is any-to-any conversion [2], where the source and target speakers are unseen during training. Here the goal is to convert to a target voice with just a few examples as a reference. Recent any-to-any methods improve naturalness and speaker similarity, but at the cost of increased complexity [3, 4]. As a result, training and inference have become more costly, making improvements difficult to evaluate and build on. We ask whether complexity is necessary for high-quality voice conversion.

To answer this question, we introduce k-nearest neighbors voice conversion (kNN-VC): a simple and robust method for any-to-any conversion. Instead of training an explicit conversion model, we simply use k-nearest neighbors regression [5]. First, we use a self-supervised speech representation model to extract feature sequences for both the source and reference utterances. Next, we convert to the target speaker by replacing each frame of the source representation with its nearest neighbor in the reference. Since certain self-supervised representations capture phonetic similarity [6], the idea is that the matched target frames would have the same content as the source. Finally, we vocode the converted features with a neural vocoder to obtain the converted speech.

Focusing on any-to-any conversion, we compare kNN-VC to several state-of-the-art systems [4, 7, 8]. Despite its simplicity,

kNN-VC matches or improves intelligibility and speaker similarity in both subjective and objective evaluations. We also run ablations to understand how target data size and vocoder training decisions affect conversion quality.

Our main contribution is to show that complexity is not necessary for any-to-any voice conversion – just applying nearest neighbors to self-supervised features gives convincing results.

## 2. Related work

Typical voice conversion systems try to separate speaker identity from content. The speaker information can then be replaced to convert to a target voice. However, learning disentangled representations is challenging [3]. Recent systems require complex information bottlenecks [4, 7, 9], normalization methods [10], or data augmentation [4], but still struggle with speaker similarity. Alternatively, some methods use text annotations as an intermediate information bottleneck to improve disentanglement [2, 8]. However, transcribing large speech corpora is costly and time-consuming. Additionally, the bottleneck imposed by text can remove prosody information which is beneficial in voice conversion applications. We therefore focus on text-free methods in this paper.

Concatenative voice conversion is another line of work related to our proposed approach [11, 12]. The idea in concatenative conversion is to stitch together units of target speech that match the source content. By constructing the output exclusively from target speech, concatenative methods ensure good speaker similarity. A key component in concatenative systems is unit selection, which generally requires parallel data [12] or time-aligned transcriptions [13]. Our method can also be seen as a unit selection approach. But instead of using human-defined units, we use features from a self-supervised speech model. Recent studies have shown that self-supervised representations linearly predict many properties of speech [14, 15]. Specifically, similar features indicate shared phonetic content [16]. This leads to our research question – are complex methods still necessary for high-quality voice conversion given these improved speech representations?

## 3. kNN-VC

To answer this question, we propose k-nearest neighbors voice conversion (kNN-VC). Figure 1 shows an overview of the method, which follows an encoder-converter-vocoder structure [17]. First, the encoder extracts self-supervised representations of the source and reference speech. Next, the converter maps each source frame to its nearest neighbor in the reference. Finally, the vocoder generates an audio waveform from the converted features. Each of these components is described below.

<sup>†</sup>Equal contribution.

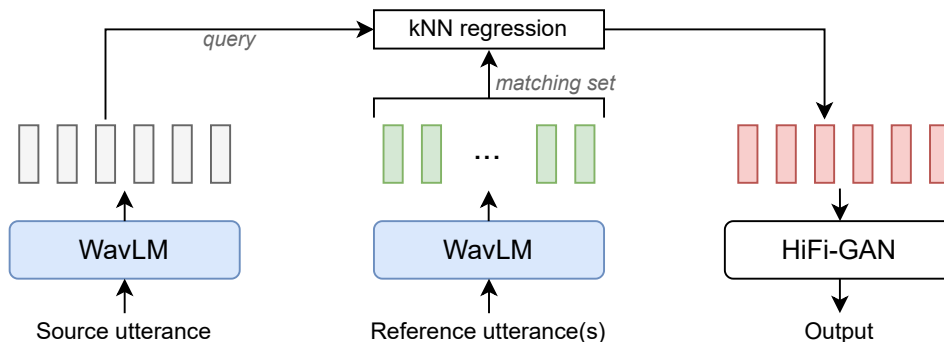


Figure 1: *kNN-VC*: The source and reference utterance(s) are encoded into self-supervised features from a pretrained WavLM model [6]. Each source feature is assigned to the mean of the  $k$  closest features from the reference. The resulting feature sequence is then vocoded using HiFi-GAN to arrive at the converted waveform output.

### 3.1. Encoder

*kNN-VC* starts by extracting the feature sequence of the source utterance, which we call the *query sequence* (Figure 1). We also extract the feature sequences of one or more utterances from the target speaker, which get shuffled together into a large pool of self-supervised feature vectors. We call this bag-of-vectors the *matching set*. The goal of this encoder is to extract representations where nearby features have similar phonetic content. Recent self-supervised models [18, 6] are good candidates since they score well on phone discrimination tests, i.e., they encode instances of the same phone closer to each other than to different phones [16]. The encoder model is never fine-tuned or trained further for *kNN-VC* – it is only used to extract features.

### 3.2. k-nearest neighbors matching

To convert to the target speaker, we apply k-nearest neighbors (kNN) regression to every vector in the query sequence. Specifically, we replace each query frame with the average of its k-nearest neighbors in the matching set. Since similar self-supervised speech features share phonetic information [6, 15], performing kNN preserves the content of the source speech while converting speaker identity. Similarly to concatenative methods, we construct the converted query directly from target features, guaranteeing good speaker similarity. The kNN regression algorithm [5] is also non-parametric and requires no training, making this method easy to implement.

### 3.3. Vocoder

The vocoder translates the converted features into an audio waveform. Instead of conditioning on spectrograms, we adapt a conventional vocoder to take self-supervised features as input. However, there is a mismatch between inputs during training and inference. For inference, we condition the vocoder on kNN outputs, i.e., the mean of features selected from the matching set. We pick these features from various points in time with different phonetic contexts, leading to inconsistencies between adjacent frames. Section 5.2 shows that these inconsistencies cause artifacts, reducing the intelligibility of converted speech.

### 3.4. Prematched vocoder training

To address this issue, we propose *prematched training*. Specifically, we reconstruct the vocoder training set using k-nearest neighbors. Using each training utterance as a query, we build

a matching set using the remaining utterances from the same speaker. Then we apply kNN regression to reconstruct the query sequence using the matching set. We train the vocoder to predict the original waveforms from these prematched features. The idea is to improve robustness by training the vocoder on data resembling what it encounters during inference. In Section 5.2, we show that prematched vocoder training improves intelligibility and speaker similarity.

## 4. Experimental setup

To determine whether our new non-parametric voice conversion method can compete with recent neural network conversion models, we run two experiments on the LibriSpeech development and test sets. The first experiment compares *kNN-VC* to state-of-the-art voice conversion systems. The second investigates the effect of target data size and prematched vocoder training. For both experiments, the source and target speakers are unseen during training. The LibriSpeech development and test clean sets consist of 40 speakers apiece, each contributing about 8 minutes of 16 kHz spoken English [19]. The diversity of speakers and audio quality in LibriSpeech make this a challenging benchmark for any-to-any voice conversion.

### 4.1. kNN-VC implementation

**Encoder:** We use the pretrained WavLM-Large encoder from [6] to extract features for the source and reference utterances. In preliminary experiments, we used features from later layers (22, 24, and the mean of the last several layers), which perform well on linear phone recognition tasks [6]. The idea was to improve nearest neighbors mapping by including more content information. However, this led to worse pitch and energy reconstruction. Recent work [15] confirms that later layers give poorer predictions of pitch, prosody, and speaker identity. Based on these observations, we found that using a layer with high correlation with speaker identification – layer 6 in WavLM-Large – was necessary for good speaker similarity and retention of the prosody information from the source utterance. So, for all the following experiments, we use features extracted from layer 6 of WavLM-Large, which produces a single vector for every 20 ms of 16 kHz audio.

**kNN regression:** For all our experiments, we set  $k = 4$  with a uniform weighting and use cosine distance to compare features. In preliminary experiments, we found that *kNN-VC* is quite

robust to a range of values around  $k = 4$ . Namely, when more reference audio is available (e.g.  $\geq 10$  mins), the conversion quality may even be improved by using larger values of  $k$  (in the order of  $k = 20$ ).

**Vocoder:** We train the HiFi-GAN V1 architecture from [20]. We modify it to accept the 1024-dimensional input vectors from WavLM and to vocode 16 kHz audio using 128-dimensional mel-spectrograms with a 10 ms hop length and 64 ms Hann window. We train the model using the same optimizer, steps and other hyperparameters as [20] on the LibriSpeech train-clean-100 dataset. We train two variants: one trained on pure WavLM-Large layer 6 features, and one trained on prematched layer 6 features (as explained in Section 3.4). With these design choices, inference with 8 minutes of reference audio on a consumer 8GB VRAM GPU is faster than real-time.

## 4.2. Baselines

We compare kNN-VC to three any-to-any voice conversion systems: VQMIVC<sup>1</sup>, FreeVC<sup>2</sup>, and YourTTS<sup>3</sup> (in text-free voice-conversion mode). VQMIVC uses vector quantization with mutual information minimization to disentangle speaker identity from content [7]. FreeVC combines a variational autoencoder with data augmentation to discard speaker information [4]. YourTTS uses text transcriptions to create an intermediate information bottleneck during training [8]. To convert to the target speaker, all three models are conditioned on speaker embeddings extracted from the reference speech using a speaker encoder. During inference, we average the speaker embeddings across all utterances from the target speaker to maximize performance. We use the publicly available, pretrained models provided by the respective authors as the baselines.

## 4.3. Evaluation metrics

We assess naturalness, intelligibility and speaker similarity in subjective and objective evaluations. We report results on the LibriSpeech test-clean subset which consists of unseen speakers not contained in any other subset of LibriSpeech.

**Objective evaluations:** To construct an evaluation set, we sample 200 utterances (5 per speaker) from the LibriSpeech test-clean set. We convert each utterance to the 39 other speakers, giving a total of 7800 outputs per model. Following previous studies [3, 7, 2], we evaluate intelligibility by calculating the word/character error rate (W/CER) of an ASR system applied to the converted speech. Lower error rates indicate better intelligibility. We use the trained Whisper-base ASR model [21] with default decoding parameters to transcribe the converted utterances.

To measure speaker similarity, we follow [22] and calculate an equal error rate (EER) using a trained speaker verification system [23]. Given an input utterance, the verification system outputs a x-vector representing the speaker identity. We compute the cosine similarity between x-vectors to determine whether two utterances are from the same speaker. First, we record similarity scores for each converted sample paired with an enrollment utterance from the target speaker. Then, we obtain an equal number of ‘ground-truth’ similarity scores by sampling target speaker utterances and computing the similarity

<sup>1</sup><https://github.com/Wendison/VQMIVC>

<sup>2</sup><https://github.com/OlaWod/FreeVC>

<sup>3</sup><https://github.com/Edresson/YourTTS>

to a different enrollment target speaker utterance. Finally, we calculate an EER over the combined set of scores, assigning a label of 1 to the ground-truth pairs and 0 to pairs containing converted speech. A higher EER indicates that it is more difficult to distinguish converted speech from genuine examples of the target speaker, i.e., higher EER corresponds to better speaker similarity, with a maximum of 50%.

**Subjective evaluations:** Using Amazon Mechanical Turk, we run subjective evaluations to measure naturalness and speaker similarity. For naturalness, we report mean opinion scores (MOS) on a 1-to-5 scale. We sample and evaluate 60 utterances for each model, adding an equal number of ground-truth examples from the LibriSpeech test-clean set. We filter out ratings from listeners who judge the ground-truth examples as very unnatural (a MOS of 1), resulting in a total of 2144 ratings from 61 listeners.

For speaker similarity, we follow the protocol from the Voice Conversion Challenge 2020 [3]: given a pair of utterances, we ask listeners to rate the similarity between the speakers on a 1-to-4 scale. For the evaluation, we sample 20 source utterance from the LibriSpeech test-clean set. We convert each utterance to three target speakers using each model, resulting in 60 examples per model. We ask raters to evaluate the similarity between the converted outputs and a genuine example of the target speaker. Again, we again filter out raters that judge the ground-truth target speaker pair as very different (a SIM of 1), resulting in a total of 1636 ratings over 208 listeners. Based on these ratings, we report a mean similarity (SIM) score for each model.

## 5. Results

We report the results of the comparative experiment in Section 5.1, followed by the ablation experiment in Section 5.2. In the first experiment, we test the voice conversion systems using the maximum available target data (about 8 minutes of audio per speaker). For kNN-VC, we use all the target data as our matching set. For the baselines, we average the speaker embeddings of each target utterance. The second experiment evaluates kNN-VC as we vary the target data size. We also investigate how prematched training affects intelligibility and speaker similarity.

### 5.1. Voice conversion

Table 1 reports intelligibility, naturalness, and speaker similarity results for each model. For kNN-VC, we use the HiFi-GAN trained on prematched data. We observe that kNN-VC achieves similar naturalness (MOS) and intelligibility (W/CER) to the best baseline, FreeVC. However, kNN-VC significantly improves

Table 1: Results measuring the intelligibility (W/CER), naturalness (MOS) and speaker similarity (EER, SIM) on the LibriSpeech test-clean subset. Subjective MOS and SIM values with 95% confidence intervals are shown.

Model	WER ↓	CER ↓	EER ↑	MOS ↑	SIM ↑
<i>Testset topline</i>	5.96	2.38	-	4.24±0.07	3.19±0.09
VQMIVC [7]	59.46	37.55	2.22	2.70±0.11	2.09±0.12
YourTTS [8]	11.93	5.51	25.32	3.53±0.09	2.57±0.12
FreeVC [4]	7.61	3.17	8.97	<b>4.07±0.07</b>	2.38±0.11
kNN-VC	<b>7.36</b>	<b>2.96</b>	<b>37.15</b>	<b>4.03±0.08</b>	<b>2.91±0.11</b>

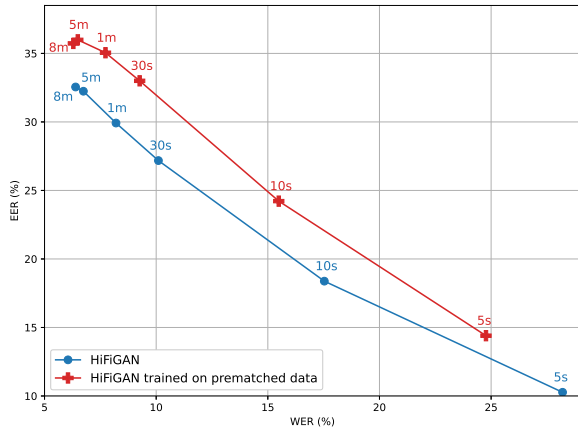


Figure 2: Intelligibility (WER, lower is better) vs speaker similarity (EER, higher is better) for varying amounts of target speaker data on the LibriSpeech dev-clean subset. The optimal result is to be in the upper-left corner.

speaker similarity (EER and SIM). Given the simplicity of kNN-VC, these results confirm our research hypothesis: increased complexity is not necessary for high-quality voice conversion.

A stronger conclusion is that our simple kNN method, combined with self-supervised speech features, leads to improved speaker similarity compared to existing models. A major benefit of kNN-VC is that it does not require an explicit speaker embedding model, in contrast to all the baselines we compare to here.

Because we do not rely on a speaker embedding model, an interesting question is: how far away can the reference utterances be from the training distribution? We leave a full investigation to future work, but include conversions to unseen languages, whispered speech, and even non-speech sounds on our demo page: <https://bshall.github.io/knn-vc>. For example, in our cross-lingual demo, we convert German source speech to a Japanese target speaker. The conversion is intelligible despite the system having only seen English during its design and training.

## 5.2. Ablation: Prematching and amount of reference data

While the main research question is answered, we would like to know how much of the improvements come from HiFi-GAN being trained on prematched data (Section 3.4), and how much the quantity of target speaker data influences intelligibility and speaker similarity. To elucidate this, Figure 2 plots WER against EER across different target speaker data sizes for both HiFi-GAN variants.

**Prematched data:** To recap from Section 3.4, prematching works by mapping each HiFi-GAN training utterance’s WavLM features to the nearest  $k$  vectors from other utterances of the same speaker (rather than just training on WavLM features directly). The higher EER and lower WER in Figure 2 shows that prematching gives a non-negligible improvement, regardless of the amount of target speaker data.

**Amount of reference data:** It is clear from Figure 2 that as amount of target speaker speech decreases, both the intelligibility and speaker similarity of the converted utterance deteriorates. This is to be expected since in the extreme case of e.g. 5 or 10 seconds of reference audio, it is unlikely that the

matching set would cover all phone or biphones necessary for smooth conversion. We see, however, that the performance gains diminishes after a certain point: using the maximum speaker data in LibriSpeech (roughly 8 minutes per speaker) gives very similar performance to only using 5 minutes of target speaker data. Again, this is likely because – beyond a certain point – there are enough vectors in the matching set to cover all possible phones and biphones present in the source utterance.

**Comparison to baselines:** Even with the plain HiFi-GAN vocoder not trained on prematched data, the scores that kNN-VC achieves in Figure 2 are still competitive with existing models in terms of intelligibility and still superior in terms of speaker similarity (comparing to the scores in Table 1). This demonstrates that, while training the vocoder on prematched data helps, using the simplest setup of a pretrained self-supervised speech model and vocoder together with kNN regression is still a powerful any-to-any voice conversion approach. However, with limited target data (less than 30s), intelligibility and target speaker similarity decrease to a point where the more complex baselines perform better (Table 1).

## 6. Conclusions

In this work, we investigated whether complex methods are necessary for any-to-any voice conversion. Using representations from a self-supervised speech model, we proposed kNN-VC: a conversion approach based on the simple kNN algorithm. We compared kNN-VC to several recent high-performing voice conversion models in a challenging any-to-any conversion scenario, and found that we achieve similar intelligibility and naturalness scores while improving speaker similarity. This confirms our main hypothesis that voice conversion can be done with simpler methods that make use of general self-supervised speech representations. We also introduced the idea of training a vocoder on prematched feature data to better approximate inference conditions, which improved performance. To better understand the limits of our approach, we looked at the effect of the amount of reference data on final voice conversion performance. We found that, with as little as five seconds of target speaker audio, we can still retain moderate intelligibility and speaker similarity.

Because the core of our approach is the vanilla kNN algorithm and our code and methods are open-source, we hope that kNN-VC can serve as a robust, easy-to-implement baseline for voice conversion. Additionally, because kNN-VC is non-parametric, we include qualitative examples of its application to cross-lingual, whispered, and even human-to-animal voice conversion – areas we hope to explore more thoroughly in future work.

## 7. Acknowledgements

All experiments were performed on Stellenbosch University’s High Performance Computing (HPC) GPU cluster. This work is supported in part by the National Research Foundation of South Africa (grant no. 120409). M. Baas is funded by the Harry Crossley Foundation.

## 8. References

- [1] S. H. Mohammadi and A. Kain, "An overview of voice conversion systems," *Speech Communication*, vol. 88, pp. 65–82, 2017.
- [2] 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.
- [3] Z. Yi, W.-C. Huang, X. Tian, J. Yamagishi, R. K. Das, T. Kinunen, Z.-H. Ling, and T. Toda, "Voice conversion challenge 2020: Intra-lingual semi-parallel and cross-lingual voice conversion," in *Joint Workshop for the Blizzard Challenge and Voice Conversion Challenge*, 2020.
- [4] J. Li, W. Tu, and L. Xiao, "FreeVC: Towards high-quality text-free one-shot voice conversion," *arXiv preprint arXiv:2210.15418*, 2022.
- [5] E. Fix, *Discriminatory analysis. Nonparametric discrimination: Consistency properties*. USAF school of Aviation Medicine, 1985, vol. 1.
- [6] 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.
- [7] 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," in *Interspeech*, 2021.
- [8] E. Casanova, J. Weber, C. D. Shulby, A. C. Junior, E. Gölge, and M. A. Ponti, "YourTTS: Towards zero-shot multi-speaker TTS and zero-shot voice conversion for everyone," in *PMLR*, 2022.
- [9] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, "AutoVC: Zero-shot voice style transfer with only autoencoder loss," in *PMLR*, 2019.
- [10] J.-c. Chou and H.-Y. Lee, "One-shot voice conversion by separating speaker and content representations with instance normalization," in *Interspeech*, 2019.
- [11] K. Fujii, J. Okawa, and K. Suigetsu, "High-individuality voice conversion based on concatenative speech synthesis," *International Journal of Electrical and Computer Engineering*, vol. 1, no. 11, pp. 1625 – 1630, 2007.
- [12] Z. Jin, A. Finkelstein, S. DiVerdi, J. Lu, and G. J. Mysore, "Cute: A concatenative method for voice conversion using exemplar-based unit selection," in *ICASSP*, 2016.
- [13] D. Sundermann, H. Hoge, A. Bonafonte, H. Ney, A. Black, and S. Narayanan, "Text-independent voice conversion based on unit selection," in *ICASSP*, 2006.
- [14] A. Pasad, J.-C. Chou, and K. Livescu, "Layer-wise analysis of a self-supervised speech representation model," in *IEEE ASRU*, 2021.
- [15] 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 *IEEE SLT*, 2023.
- [16] E. Dunbar, N. Hamilakis, and E. Dupoux, "Self-supervised language learning from raw audio: Lessons from the zero resource speech challenge," *Journal of Selected Topics in Signal Processing*, 2022.
- [17] B. Sisman, J. Yamagishi, S. King, and H. Li, "An overview of voice conversion and its challenges: From statistical modeling to deep learning," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 132–157, 2021.
- [18] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS*, 2020.
- [19] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR corpus based on public domain audio books," in *ICASSP*, 2015.
- [20] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *NeurIPS*, 2020.
- [21] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *arXiv preprint arXiv:2212.04356*, 2022.
- [22] B. van Niekerk, M.-A. Carbonneau, J. Zaïdi, M. Baas, H. Seuté, and H. Kamper, "A comparison of discrete and soft speech units for improved voice conversion," in *ICASSP*, 2022.
- [23] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-Vectors: Robust dnn embeddings for speaker recognition," in *ICASSP*, 2018.