



Discrete Tokens Exhibit Interlanguage Speech Intelligibility Benefit: an Analytical Study Towards Accent-robust ASR Only with Native Speech Data

Kentaro Onda^{1,2}, Keisuke Imoto³, Satoru Fukayama², Daisuke Saito¹, Nobuaki Minematsu¹

¹The University of Tokyo, Japan

²National Institute of Advanced Industrial Science and Technology (AIST), Japan

³Kyoto University, Japan

{ondakentaro, mine}@gavo.t.u-tokyo.ac.jp

Abstract

In this study, we gained insight that contributes to achieving accent-robust ASR using only native speech data. In human perception of non-native speech, the phenomenon known as “interlanguage speech intelligibility benefit” (ISIB) is observed, where non-native listeners who share the native language with the speaker understand the speech better compared even to native listeners. Based on the idea that discrete tokens extracted from self-supervised learning (SSL) models represent the human perception of speech, we conducted an analytical study on the robustness of discrete token-based ASR to non-native speech, varying the language used for training the tokenization, which is viewed as a technical implementation of ISIB. The results showed that ISIB actually occurred in the discrete token-based ASR. Since our approach relies only on native speech data to simulate the behavior of human perception, it is expected to be applicable to a wide range of accents for which speech data is scarce.

Index Terms: discrete tokens, self-supervised learning, foreign accent, interlanguage speech intelligibility benefit, automatic speech recognition

1. Introduction

Automatic speech recognition (ASR) is widely used in applications such as voice-controlled devices and transcription. ASR systems are required to be robust against audio variations caused by factors such as speakers and recording environments. However, it is well known that ASR performance tends to degrade when recognition is performed on accented speech by non-native speakers [1]. In today’s increasingly globalized world, communication between people with different native languages is becoming more common. Furthermore, in recent language education, the “intelligibility principle” has become dominant, which emphasizes that learners’ pronunciation does not need to be native-like as long as it is intelligible enough [2, 3, 4]. As a result, people are more likely to speak with their own accents, and the demand for ASR systems capable of recognizing foreign-accented speech has grown even further.

To achieve accent-robust ASR, various methods such as data augmentation and accent adaptation have been proposed [5, 6, 7, 8]. However, these methods generally require accented speech data for training, limiting their applicability only to a few types of accents for which such corpora are available. They are mostly of “X-accented English” [9], while there are much more diverse combinations of “X-accented Y” spoken in the real world. On the other hand, native speech data, both for language X and Y, are relatively easier to obtain for various languages. If these data are used effectively to achieve accent-robust ASR, such approach could be applicable to a wide range

of language accents, including those for which accented speech data are scarce or even unavailable.

In [10], a method was proposed to synthesize foreign-accented speech using only native speech data. This approach assumes the resynthesis of speech using discrete tokens extracted from self-supervised learning (SSL) models as a model of human speech reproduction, where speech is perceived as discrete symbols, which are then vocalized. By resynthesizing the input speech with a model trained on a different language, the model can generate speech that sounds like the reproduction of the input speech by a non-native speaker. This method is based on the hypothesis that the process of discretization simulates how humans with different native languages perceive speech. For example, discrete tokens trained on Japanese speech are assumed to represent the results of the speech perception experienced by native Japanese speakers.

In the context of human perception of non-native speech, a phenomenon known as “interlanguage speech intelligibility benefit” (ISIB) is often observed [11, 12, 13]. This refers to the tendency where non-native speech is more intelligible to listeners who share the same native language with the speaker than to native listeners of the spoken language. For example, Japanese listeners tend to have an advantage in understanding Japanese-accented English compared to native English listeners.

In this study, as a first step towards accent-robust ASR using only native speech data, we investigated whether ISIB can be observed in discrete token-based ASR. Specifically, we examined whether the recognition performance for foreign-accented speech can be improved by using discrete tokens trained on the speaker’s native language. This could also validate the hypothesis that discrete tokens indeed represent the human perception of speech. Unlike synthesis task performed in [10], where acoustic information is added to the tokens during vocalization process, the output of ASR consists purely of linguistic information. This makes it possible to conduct the validation in a more interpretable way.

Our scientific contributions are summarized as follows:

- We conducted an analytical study on the robustness of discrete token-based ASR to foreign-accented speech, changing the language used for training the tokenization as a technical implementation of ISIB.
- We confirmed the occurrence of ISIB in the discrete token-based ASR, which provides additional support for the hypothesis that discrete tokens represent the human perception of speech.
- We suggested the potential of our approach, using discrete tokens to simulate ISIB, for achieving accent-robust ASR using only native speech data.




Figure 1: The reproduction of native speech of language Y by (a) native speakers of language Y, (b) native speakers of language X: the output speech of (b) is expected to be “X-accented Y”.




Figure 2: The recognition of “X-accented Y” speech by (a) native speakers of language Y, (b) native speakers of language X: (b) is expected to be more accurate through the ISIB effect.

2. Related studies

2.1. Generative Spoken Language Model

Generative Spoken Language Model (GSLM) [14] is a language model trained only with speech data without any text. It was proposed under the concept of “textless NLP,” based on the idea that text is not always necessary for human language acquisition. In GSLM, “pseudo-text” is used as a substitute for text, which is learned through k-means clustering of features extracted from SSL models such as HuBERT [15] and wav2vec2.0 [16]. This pseudo-text, also referred to as “units” or “discrete tokens,” has been explored for various applications across different tasks [17]. Our experiment also utilizes this discrete tokens as intermediate representations for ASR.

2.2. Simulation of foreign accentuation only with native speech data

In [10], a method for synthesizing foreign-accented speech using only native speech data was proposed. This method is based on the idea that the discrete tokens in GSLM represents the human perception of speech. Foreign accents occur when speech in a foreign language is perceived within the framework of the listener’s native language. When the listener repeats it, the perceived speech is vocalized in a way also influenced by the native language [18]. This human process of speech reproduction was modeled as speech resynthesis using discrete tokens as intermediate representations, as shown in Figure 1. By training the tokenization and its vocalization with native speech data of language X, which is different from the input speech of language Y, it is possible to generate “X-accented Y” speech. Our study extends this idea of regarding discrete tokenization as a model of human perception to ASR for foreign-accented speech.

2.3. Discrete token-based ASR

In [19], ASR using discrete tokens as intermediate representations was proposed. Discrete tokens, which represent each frame as a single integer, offer high efficiency in terms of computation and memory usage. In addition to this, this approach achieves higher recognition performance compared to conventional methods with acoustic features such as log Mel-

filterbanks. Subsequently, research on discrete token-based ASR has been actively conducted [20, 21, 22, 23]. However, these studies generally use the same dataset, thus same language, for training the discrete tokens and the ASR model. The evaluation is also conducted on the same in-domain data. In this study, we focus on the use of discrete tokens trained on a different language and examine the robustness of ASR against foreign-accented speech not included in the training data.

3. Empirical analyses

Figure 2 illustrates an overview of discrete token-based ASR for foreign-accented speech, aiming to simulate ISIB. Based on the concept of [10], shown in the section 2.2, discrete tokens are regarded as results of human perception. When recognizing “X-accented Y” speech, by using discrete tokens trained on language X, the native language of the speaker, the perceptive behaviors of listeners with the same native language X are simulated. Then the recognition performance is expected to be improved through ISIB. In this approach, ASR is trained only on native speech data of language Y, but the data are all converted into discrete tokens of language X. The tokenization is also trained only on native speech data of language X. Therefore, this approach aims to enhance recognition accuracy for “X-accented Y” speech using only native speech data of languages X and Y. This makes the method particularly effective in scenarios where accented speech data is scarce.

3.1. Experimental setup

The experiments were conducted using ESPnet [24], with the joint CTC/attention-based encoder-decoder ASR model [25] following the configurations described in [19, 20]. For input token sequences, subword modeling with byte pair encoding was not performed to facilitate comparison across different cluster sizes, and only deduplication (removing consecutive identical tokens) was applied. We employed HuBERT-base [15] as the SSL model for feature extraction, following [10, 14]. In all experiments, the ASR training was done for 50 epochs, and the model averaged over the top 3 epochs based on the validation set was used for evaluation.

First, experiments on Japanese-accented English were conducted to be described in sections from 3.2 to 3.4. In this scenario, we compared the recognition performance for Japanese-accented English under two conditions: when discretization via k-means clustering was trained on English or Japanese. For training k-means clusters, about 30 hours of native speech data were used for each language. A random subset of the train-clean-100 of LibriSpeech [26] was used for English, and JVS [27] was used for Japanese. For evaluation, the ERJ corpus [28] was used. This corpus consists of English speech read by both Japanese learners and native speakers of American English. In our experiments, 460 phoneme-balanced sentences were used. For Japanese speakers’ speech, pronunciation assessment annotations are available in the corpus, and we used the segmental scores for evaluation as a measure of accent strength.

Lastly, in 3.5, experiments were extended to cover a wider range of accents. Detailed settings are described in that section.

3.2. LibriSpeech100 results

We first trained and evaluated ASR models using LibriSpeech100 dataset. In this experiment, comparisons were made not only with the language used for k-means training but also the language for SSL pretraining and the cluster sizes. We

Table 1: ASR performance comparison: WER[%] (\downarrow) (trained on LibriSpeech100)

cluster size	SSL	km	LibriSpeech	ERJ	
			test-{clean, other}	AE	JE_all/JE_w10
100	en	en	10.1/24.0	30.0	91.1/101.9
		jp	15.2/30.1	40.4	94.8/104.7
	jp	en	17.9/36.7	45.1	96.3/105.8
		jp	21.7/41.4	51.7	100.4/108.9
500	en	en	8.0/19.2	23.6	83.8/95.9
		jp	9.1/21.1	27.4	82.7/95.2
	jp	en	13.0/29.7	34.4	88.9/100.3
		jp	14.7/32.0	39.1	87.4/97.0
2000	en	en	7.4/17.4	22.1	80.4/93.4
		jp	8.1/18.9	23.9	79.5/92.0
	jp	en	10.7/26.2	30.8	86.5/98.0
		jp	13.1/29.9	37.1	87.4/98.0

compared two pretrained HuBERT-base models: one trained on English¹ and the other one trained on Japanese². In both cases, features were extracted from the 12th (last) layer of the model, and three cluster sizes were evaluated: 100, 500, and 2000.

The results are shown in Table 1, where four conditions of the language for SSL pre-training and k-means training (SSL, km) were tested for each cluster size. The evaluation data includes the test sets of LibriSpeech, as well as the ERJ corpus. For ERJ, American English (AE) and Japanese English (JE) were evaluated separately. For JE, the results were shown for all utterances (JE_all) and for the subset from the 10 worst speakers based on the segmental score annotations (JE_w10), who were supposed to have the strongest Japanese accent.

As for comparison between evaluation sets, in all cases, in-domain LibriSpeech showed the highest accuracy, followed by ERJ (AE), the out-of-domain native English data, and then ERJ (JE), the out-of-domain Japanese-accented English. Among ERJ (JE), the JE_w10 subset showed lower accuracy than JE_all. Increasing the number of clusters improved performance for all test sets. This trend was also discussed in a previous study [19], and now it was confirmed that the same trend was observed even with out-of-domain and accented speech data. When using an SSL model pre-trained in Japanese, the accuracy deteriorated compared to the English-pretrained model in all cases. For SSL features, it is considered effective to use models trained on the target language (in this case, English) regardless of accent.

When using the SSL model pre-trained in English, for native English (LibriSpeech and ERJ (AE)), the accuracy was higher when the k-means clustering was trained on English data. However, for Japanese-accented English (ERJ (JE)), higher accuracy was observed when k-means was trained on Japanese, when the cluster size was 500 or 2000. This suggests that ISIB, where Japanese listeners have an advantage in understanding Japanese-accented English, may have been replicated through the use of discrete tokens. This trend was not observed when the cluster size was 100, probably because the number of clusters was too small for effective speech recognition. In fact, when the cluster size was 100, recognition accuracy even for native English speech also deteriorated significantly compared to when the cluster sizes were 500 or 2000.

3.3. LibriSpeech960 results

Next, the experiment was extended to the LibriSpeech960 to validate the occurrence of ISIB with a larger dataset. The num-

¹<https://huggingface.co/facebook/hubert-base-ls960>

²<https://huggingface.co/rinna/japanese-hubert-base>

Table 2: ASR performance comparison: WER[%] (\downarrow) (trained on LibriSpeech960)

cluster size	SSL (layer)	km	LibriSpeech	ERJ	
			test-{clean, other}	AE	JE_all/JE_w10
2000	en (12)	en	4.0/9.2	13.4	55.7/ 70.8
		jp	3.8/9.2	14.3	53.3/68.0
		Δ (en-jp)	0.2/0.0	-0.9	2.4/2.8
en (9)	en	en	3.5/8.8	12.3	59.9/74.9
		jp	4.1/9.9	14.9	55.2/71.2
		Δ (en-jp)	-0.6/-1.1	-2.6	4.7/3.7
en (6)	en	en	3.7/9.4	13.5	56.9/71.3
		jp	4.8/11.0	16.8	54.0/67.2
		Δ (en-jp)	-1.1/-1.6	-3.3	2.9/4.1

Table 3: An example of recognition results

Recognition result of IWA/M02/S2_002	
GT	any contributions will be greatly appreciated
en (9)	any contribution* will be greatly appliciated
jp (9)	any contribution* will be greatly appreciated

ber of clusters was set to 2000, and the HuBERT-base pretrained in English was used based on the results of the previous section. In addition to the last layer, the cases with the 9th and 6th layers were also evaluated.

The results are shown in Table 2. In all cases, for the out-of-domain ERJ data, AE were better recognized when k-means clustering was done on English speech, while JE showed better recognition accuracy when k-means clusters were trained on Japanese. This clearly exhibits the occurrence of ISIB. Among different SSL layers, there was a general trend where the difference between languages for k-means training (Δ (en-jp)) became smaller as the layers approached the last layer. This is likely because upper layers have less acoustic information and more linguistic information [29], allowing subtle pronunciation differences caused by accents to be less captured.

Table 3 shows an example of recognition results of Japanese accented English, where /r/ was mispronounced like /l/ (a common error among Japanese speakers [30, 31]). The model trained with English k-means misrecognized the word ‘‘appreciated’’ as ‘‘appliciated,’’ while the model trained with Japanese k-means correctly recognized the word. This suggests that clusters trained in English may distinguish the sounds of /r/ and /l/ more clearly, while those trained in Japanese are more tolerant of the difference, making them more robust to Japanese accents.

3.4. The evaluation of clusters

The evaluation of the clustering quality is typically conducted by comparing the discrete token sequences with the ground truth phoneme sequences, such as phone normalized mutual information (PNMI) [14, 15, 19]. However, for non-native speech, it is difficult to define accurate phoneme labels. Therefore, in this study, we explore the following two metrics that do not rely on the phonemic transcripts.

1) Quantization Error (QE): The squared Euclidean distance between the original SSL features and the centroids of the clusters to which the features are assigned. A smaller value indicates that the discretization is more suitable for the given data.

2) Mean Token Error Rate (MTER): The average of token error rate (TER) computed across all possible pairs of multiple utterances of the same content. TER is defined as the edit distance between two discrete token sequences generated from two utterances of the same content. A lower MTER suggests that the tokenization is more robust to non-linguistic variations, including accents. Unlike word error rate (WER) or character

Table 4: Cluster quality evaluation with ERJ corpus

cluster size	SSL (layer)	km	QE (↓)		MTER [%] (↓)	
			AE	JE_all/JE_w10	AE	JE_all/JE_w10
100	en (12)	en	79.1	83.2/84.3	39.4	79.8/94.7
		jp	94.9	85.0/84.8	50.8	83.2/90.1
500	en (12)	en	56.9	63.7/65.0	53.7	98.5/114.1
		jp	78.8	67.7/67.3	58.6	101.5/112.3
2000	en (12)	en	47.1	54.9/56.3	66.7	108.7/121.1
		jp	66.8	57.6/57.5	70.4	115.4/131.8
	en (9)	en	30.7	38.0/39.8	53.7	103.2/120.7
		jp	43.5	39.1/39.4	64.5	106.8/121.9
	en (6)	en	30.6	37.1/38.7	64.3	107.9/123.9
		jp	41.4	38.9/40.1	73.5	112.3/125.2

error rate (CER), MTER is calculated using multiple utterances as reference data for each sentence, as ground truth token sequences cannot be defined. In the ERJ corpus, the same sentence is read by 10 native speakers and 24 Japanese learners. We used all the native speakers’ utterances as reference data and calculated TER for each utterance, then averaged the results.

The results are shown in Table 4. The evaluation was conducted using the ERJ corpus, and only the results for k-means models using English-pretrained SSL are shown.

Regarding QE, in most cases, the value was lower when k-means clustering was trained on English than on Japanese. However, while when clusters are trained on English, AE showed lower QE than JE, the opposite trend was observed when k-means was trained with Japanese data. As the proficiency of the speaker decreases as AE → JE_all → JE_w10, the difference between the results of English and Japanese k-means becomes smaller. When using the 9th layer, QE with Japanese k-means got lower than with English k-means for JE_w10, indicating that Japanese k-means gives more adequate units for speech with stronger Japanese accents.

For MTER as well, k-means learned on English showed lower values in most cases. However, when the cluster size was 100 or 500, the model trained on Japanese data showed lower MTER for JE_w10. This is likely because speakers with lower proficiency tend to have stronger accents, and Japanese sounds are more often substituted in the JE_w10 utterances. As a result, the distribution in the SSL feature space tends to fit better with the cluster centroids learned with Japanese data. The reason why the same trend was not observed with the cluster size of 2000 is likely because as the number of clusters increased, the case became more frequent where the same sentence spoken with different accents was assigned to different clusters. In the comparison between layers, the 9th layer showed lower MTER for AE than other layers. We could say that this is related to the difference of ASR performances among layers.

3.5. Do mismatched ISIBs occur?

Several studies reported the existence of “mismatched” ISIB, where non-native listeners have an advantage in understanding non-native speech even when the native languages of the speaker and listener are different [11, 32]. While this is not always true for human listeners [33, 34], we investigated whether such a phenomenon could be seen in discrete token-based ASR.

In this section, we examined more diverse cases where X-accented Y speech was recognized using discrete tokens trained on either language X (“matched”) or language Z (“mismatched”). We compared ASR performance on accented speech using k-means trained not only on English and Japanese but also on Chinese and Spanish. For k-means training in Chinese and Spanish, we used randomly selected 30-hour subsets from the AISHELL [35] and MLS (Spanish) [36], respectively.

Table 5: ASR results for more diverse accents. Cluster size: 2000, SSL: en (9). WER[%] (↓) (trained on LibriSpeech960)

km	ERJ		L2-ARCTIC					
	AE	JE_all/JE_w10	Chinese	Spanish	Arabic	Hindi	Korean	Vietnamese
en	12.3	<u>59.9/74.9</u>	29.8	23.4	<u>23.5</u>	<u>20.0</u>	19.6	35.9
jp	14.9	55.2/71.2	<u>30.1</u>	23.3	23.4	19.8	19.2	<u>36.9</u>
ch	13.7	57.3/71.6	28.6	22.4	22.4	19.6	18.6	34.9
sp	14.4	57.1/71.4	29.3	21.8	22.1	19.3	18.9	35.3

ASR training was done with LibriSpeech960, and the cluster size was set to 2000. The 9th layer output of HuBERT-base was employed since the layer showed the best performance for native English in the previous section. For evaluation, in addition to ERJ, the L2-ARCTIC corpus [37] was used. This corpus consists of English speech read by non-native speakers from 6 different native languages, including Chinese and Spanish.

The results are shown in Table 5. For each test set, the best and worst results are shown in bold and underlined, respectively. For all the “matched” case, as shown in the diagonal elements of the first four columns, the recognition accuracy was highest when k-means clustering was trained on the same language as the speaker’s native language. This suggests that the “matched” ISIB was also observed in Chinese and Spanish-accented English. As for the “mismatched” cases, when k-means was trained on Chinese or Spanish, the recognition accuracy for all foreign accents was higher than when trained on English. In most cases, k-means trained on English showed the worst performance. This suggests that the “mismatched” ISIB occurs in the discrete token-based ASR. We may say that this approach is effective even when the speakers’ native language is unknown if only they are non-native speakers.

For Arabic and Hindi-accented English, the recognition accuracy was the highest when k-means clustering was trained on Spanish. For Korean and Vietnamese-accented English, when using Chinese-trained k-means showed the best performance. This suggests that the “mismatched” ISIB on the discrete token-based ASR is more effective when the accents of the speaker and the listener are closer. In fact, a study on accent classification showed some confusion between Arabic and Spanish, and between Vietnamese and Chinese accents [38]. This indicates that this method can be applied even when the speaker’s native language is low-resource and it is difficult to obtain enough speech data for k-means training. By selecting a more high-resource language the accent of which is similar to the speaker’s accent, the recognition performance may be further improved.

4. Conclusions

In this study, we conducted comparative experiments on discrete token-based ASR for foreign-accented speech varying the language used for k-means clustering. The results showed that training k-means on the speaker’s native language led to the best recognition performance among the tested languages, including the spoken language. This shows the occurrence of ISIB in discrete token-based ASR, and supports the idea suggested in [10] more strongly that discrete tokens simulate the human perception of speech. Since our approach of replicating ISIB relies only on native speech data, it can be applied even if the accented speech data is unavailable. The fact that “mismatched” ISIB was also observed in discrete token-based ASR shows further applicability of this approach, even when the speaker’s native language is unknown or low-resource. Future work includes exploring more advanced simulations of speech perception by multilingual listeners and improving recognition performance through them while using only native speech data.

5. Acknowledgements

This work was supported by AIST KAKUSEI project (FY2024).

6. References

- [1] A. Koenecke, A. Nam, E. Lake, J. Nudell, M. Quartey, Z. Mengesha, C. Troups, J. R. Rickford, D. Jurafsky, and S. Goel, "Racial disparities in automated speech recognition," *PNAS*, vol. 117, no. 14, pp. 7684–7689, 2020.
- [2] M. J. Munro and T. M. Derwing, "Foreign accent, comprehensibility, and intelligibility in the speech of second language learners," *Language Learning*, vol. 45, no. 1, pp. 73–97, 1995.
- [3] J. M. Murphy, "Intelligible, comprehensible, non-native models in ESL/EFL pronunciation teaching," *System*, vol. 42, pp. 258–269, 2014.
- [4] J. Levis, "Revisiting the intelligibility and nativeness principles," *Journal of Second Language Pronunciation*, vol. 6, no. 3, pp. 310–328, 2020.
- [5] T. Fukuda, R. Fernandez, A. Rosenberg, S. Thomas, B. Ramabhadran, A. Sorin, and G. Kurata, "Data augmentation improves recognition of foreign accented speech," in *Interspeech 2018*, 2018, pp. 2409–2413.
- [6] P. Klumpp, P. Chitkara, L. Sari, P. Serai, J. Wu, I.-E. Veliche, R. Huang, and Q. He, "Synthetic cross-accent data augmentation for automatic speech recognition," in *arXiv*, 2023, p. 2303.00802.
- [7] D. Prabhu, P. Jyothi, S. Ganapathy, and V. Unni, "Accented speech recognition with accent-specific codebooks," in *EMNLP 2023*, 2023, pp. 7175–7188.
- [8] D. Prabhu, A. Gupta, O. Nitsure, P. Jyothi, and S. Ganapathy, "Improving self-supervised pre-training using accent-specific codebooks," in *Interspeech 2024*, 2024, pp. 2310–2314.
- [9] A. Hinsvark, N. Delworth, M. D. Rio, Q. McNamara *et al.*, "Accented speech recognition: A survey," in *arXiv*, 2021, p. 2104.10747.
- [10] K. Onda, J. Park, N. Minematsu, and D. Saito, "A pilot study of GSLM-based simulation of foreign accentuation only using native speech corpora," in *Interspeech 2024*, 2024, pp. 3600–3604.
- [11] T. Bent and A. R. Bradlow, "The interlanguage speech intelligibility benefit," *JASA*, vol. 114, no. 3, pp. 1600–1610, 2003.
- [12] L. Harding, "Accent, listening assessment and the potential for a shared-L1 advantage: A DIF perspective," *Language Testing*, vol. 29, no. 2, pp. 163–180, 2012.
- [13] X. Xie and C. A. Fowler, "Listening with a foreign-accent: The interlanguage speech intelligibility benefit in Mandarin speakers of English," *Journal of Phonetics*, vol. 41, no. 5, pp. 369–378, 2013.
- [14] K. Lakhotia, E. Kharitonov, W.-N. Hsu, Y. Adi *et al.*, "On generative spoken language modeling from raw audio," *TACL*, vol. 9, pp. 1336–1354, 2021.
- [15] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *T-ASLP*, vol. 29, pp. 3451–3460, 2021.
- [16] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS 2020*, vol. 33, 2020, pp. 12 449–12 460.
- [17] X. Chang, J. Shi, J. Tian, Y. Wu, Y. Tang, Y. Wu, S. Watanabe, Y. Adi, X. Chen, and Q. Jin, "The Interspeech 2024 challenge on speech processing using discrete units," in *Interspeech 2024*, 2024, pp. 2559–2563.
- [18] J. E. Flege, "Second language speech learning: Theory, findings, and problems," *Speech perception and linguistic experience: Issues in cross-language research*, vol. 92, pp. 233–277, 1995.
- [19] X. Chang, B. Yan, Y. Fujita, T. Maekaku, and S. Watanabe, "Exploration of efficient end-to-end ASR using discretized input from self-supervised learning," in *Interspeech 2023*, 2023, pp. 1399–1403.
- [20] X. Chang, B. Yan, K. Choi, J.-W. Jung *et al.*, "Exploring speech recognition, translation, and understanding with discrete speech units: A comparative study," in *ICASSP 2024*, 2024, pp. 11 481–11 485.
- [21] Y. Yang, F. Shen, C. Du, Z. Ma, K. Yu, D. Povey, and X. Chen, "Towards universal speech discrete tokens: A case study for ASR and TTS," in *ICASSP 2024*, 2024, pp. 10 401–10 405.
- [22] J. Shi, X. Ma, H. Inaguma, A. Sun, and S. Watanabe, "MMM: Multi-layer multi-residual multi-stream discrete speech representation from self-supervised learning model," in *Interspeech 2024*, 2024, pp. 2569–2573.
- [23] P. Mousavi, J. Duret, S. Zaiem, L. Della Libera, A. Ploujnikov, C. Subakan, and M. Ravanelli, "How should we extract discrete audio tokens from self-supervised models?" in *Interspeech 2024*, 2024, pp. 2554–2558.
- [24] S. Watanabe, T. Hori, S. Karita, T. Hayashi *et al.*, "ESPnet: End-to-end speech processing toolkit," in *Interspeech 2018*, 2018, pp. 2207–2211.
- [25] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *ICASSP 2017*, 2017, pp. 4835–4839.
- [26] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *ICASSP 2015*, 2015, pp. 5206–5210.
- [27] S. Takamichi, K. Mitsui, Y. Saito, T. Koriyama, N. Tanji, and H. Saruwatari, "JVS corpus: free Japanese multi-speaker voice corpus," 2019.
- [28] N. Minematsu, Y. Tomiyama, K. Yoshimoto, K. Shimizu, S. Nakagawa, M. Dantsuji, and S. Makino, "Development of English speech database read by Japanese to support CALL research," in *ICA 2004*, 2004, pp. 557–560.
- [29] P. Kumar, V. N. Sukhadia, and S. Umesh, "Investigation of robustness of hubert features from different layers to domain, accent and language variations," in *ICASSP 2022*, 2022, pp. 6887–6891.
- [30] A. R. Bradlow, D. B. Pisoni, R. Akahane-Yamada, and Y. Tohkura, "Training Japanese listeners to identify English/r/and/l: Iv. some effects of perceptual learning on speech production," *JASA*, vol. 101, no. 4, pp. 2299–2310, 1997.
- [31] R. M. McKenzie, "The sociolinguistics of variety identification and categorisation: free classification of varieties of spoken English amongst non-linguist listeners," *Language Awareness*, vol. 24, no. 2, pp. 150–168, 2015.
- [32] J. Yuan, Y. Jiang, and Z. Song, "Perception of foreign accent in spontaneous L2 English speech," in *Speech Prosody 2010*, 2010.
- [33] H. Wang and V. J. Van Heuven, "The interlanguage speech intelligibility benefit as bias toward native-language phonology," *i-Perception*, vol. 6, no. 6, p. 2041669515613661, 2015.
- [34] R. M. Stibbard and J.-I. Lee, "Evidence against the mismatched interlanguage speech intelligibility benefit hypothesis," *JASA*, vol. 120, no. 1, pp. 433–442, 2006.
- [35] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "AISHELL-1: An open-source Mandarin speech corpus and a speech recognition baseline," in *O-COCOSDA 2017*, 2017, pp. 1–5.
- [36] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "MLS: A large-scale multilingual dataset for speech research," in *Interspeech 2020*, 2020, pp. 2757–2761.
- [37] G. Zhao, S. Sonsaat, A. Silpachai, I. Lucic, E. Chukharev-Hudilainen, J. Levis, and R. Gutierrez-Osuna, "L2-ARCTIC: A non-native English speech corpus," in *Interspeech 2018*, 2018, pp. 2783–2787.
- [38] M. Yang, R. C. M. C. Shekar, O. Kang, and J. H. L. Hansen, "What can an accent identifier learn? probing phonetic and prosodic information in a wav2vec2-based accent identification model," in *Interspeech 2023*, 2023, pp. 1923–1927.