



How Generative Spoken Language Modeling Encodes Noisy Speech: Investigation from Phonetics to Syntactics

Joonyong Park¹, Shinnosuke Takamichi¹, Tomohiko Nakamura¹,
Kentaro Seki¹, Detai Xin¹, Hiroshi Saruwatari¹

¹The University of Tokyo, Japan

joonyong-park@g.ecc.u-tokyo.ac.jp, shinnosuke_takamichi@ipc.i.u-tokyo.ac.jp

Abstract

We examine the speech modeling potential of generative spoken language modeling (GSLM), which involves using learned symbols derived from data rather than phonemes for speech analysis and synthesis. Since GSLM facilitates textless spoken language processing, exploring its effectiveness is critical for paving the way for novel paradigms in spoken-language processing. This paper presents the findings of GSLM’s encoding and decoding effectiveness at the spoken-language and speech levels. Through speech resynthesis experiments, we revealed that resynthesis errors occur at the levels ranging from phonology to syntactics and GSLM frequently resynthesizes natural but content-altered speech.

Index Terms: speech synthesis, noise reduction, speech correction, self-supervised learning

1. Introduction

The field of natural language processing has made significant advancements [1], particularly in the area of language models. Studies have shown that language models can naturally learn multiple tasks even without supervision, which can enable the model to comprehend, answer questions, summarize, and translate text [2–4]. Despite the success in text-level language modeling, it does not handle spoken-language expressions, which is another essential component of speech. To resolve this problem, investigations on speech-language models through the use of extensive unlabeled speech data have been conducted [5–10].

The most promising approach to this problem is generative spoken language modeling (GSLM) [11], which involves learning the acoustic and spoken linguistic features only from speech audio. As shown in Fig. 1(a), GSLM involves analyzing and synthesizing speech audio through a sequence of discrete symbols. These symbols can be viewed as pseudo-phones, though they are not exactly equivalent to phonemes or phones as those defined in phonetics [11–13]. Since these symbols can be determined through training, GSLM does not require transcriptions associated with speech data, i.e., it enables text-free spoken language processing [14–16]. It also potentially expands the available speech data beyond carefully designed and annotated laboratory data to real-world data.

However, real-world data usually contain noise and sometimes include speech of non-target speakers. This may affect the encoding performance in GSLM [17], posing a challenge for GSLM in being robust to noise contamination. This motivated us to investigate what information GSLM encodes from noisy speech and identify which levels of spoken language, ranging from phonetic to syntactic levels, are affected by noise.

T. Nakamura is currently with National Institute of Advanced Industrial Science and Technology (AIST), Japan.

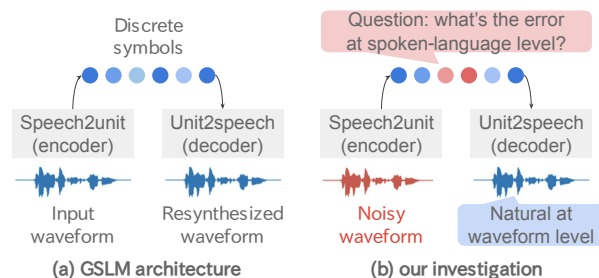


Figure 1: *GSLM-model architecture and our investigation.*

We investigated the encoding and decoding effectiveness of GSLM in a noisy environment (see Fig. 1(b)). To do this, we conducted experiments on speech resynthesis from noisy speech using the GSLM model trained with speech without noise. The experimental results were analyzed at the spoken-language and speech levels.

2. Generative spoken language modeling

This section outlines the architecture of the GSLM model [11]. Figure 1(a) shows a schematic illustration of the GSLM-model architecture. The GSLM model analyzes and synthesizes an audio waveform of speech through discrete symbols (units) instead of phoneme symbols. The primary components of the GSLM model are speech2unit and unit2speech modules. The speech2unit module first converts the audio waveform into framewise features using the encoder of a self-supervised speech representation model (e.g., contrastive predictive coding [5], wav2vec2.0 [8], and HuBERT [9]). It then quantizes the features with a predetermined codebook to obtain the discrete symbol sequence. The codebook is obtained by applying k -means clustering to the framewise features of the training data. The unit2speech module decodes the discrete symbol sequence into an audio waveform. For this module, we can use the network architectures used in conventional text-to-speech synthesis models and neural vocoder models (e.g., Tacotron2 [18]).

Another important module of the GSLM model is a unit-based language model (uLM), which is omitted in Figure 1. It serves as a generative language model in the unit domain and enables speech generation and continuation without transcription. However, the focus of this study was on speech resynthesis, which does not require the uLM. Thus, we analyze the speech2unit and unit2speech modules.

3. Experimental analysis on noisy speech

3.1. Experimental setup

We analyzed the robustness of GSLM to noise contamination through speech resynthesis experiments. We chose Hu-

Table 1: *Categorization of noise types of DEMAND*

Category	Noise type of DEMAND
L-BAB	DKITCHIN, DWASHING, NFIELD, NRIVER, OHALLWAY, OOFFICE, TCAR
M-BAB	DLIVING, NPARK, PSTATION, SPSQUARE, STRAFFIC
H-BAB	OMEETING, PCAFETER, PRESTO, TBUS, TMETRO

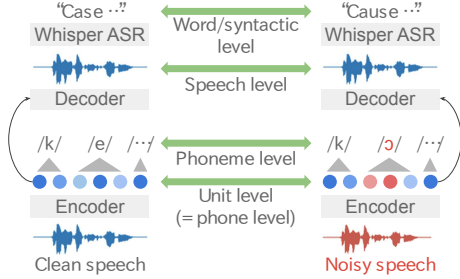


Figure 2: *Levels investigated in experiments.*

BERT [19] with the codebook of 200 classes as the speech2unit module and Tacotron2 [18] as the unit2speech module, which achieved the highest performance in word error rate (WER) and mean opinion score of synthesized speech in [11]. The GSLM model, in which the encoder is trained on the LibriSpeech corpus [20] and decoder is trained on the LJSpeech corpus [21] is available via the fairseq toolkit [22].

The noisy speech signals were created by adding noise signals to clean speech signals with varying signal-to-noise ratios (SNRs) from 0 to 15 dB with an interval of 5 dB. We used the dev-clean set of the LibriSpeech corpus (2704 utterances) [20] as the clean-speech signals. The noise signals were drawn from the diverse environments multichannel acoustic noise database (DEMAND) [23], which contains noise signals recorded with a 16-channel microphone array in 17 noisy environments. We randomly chose one of the channels for each noise signal. The number of the test dataset was 183,872 in total.

To examine the effect of background speech, we classified the noise types into three categories along with the amount of speech contained in the noise. Table 1 shows the categorization of the noise types. See [23] for details of these noise types. **L-BAB** is the collection of the noise types that contain little or no speech, i.e., low babble noise. **M-BAB** contains the noise types including speech in less than half the signal length. **H-BAB** is the collection of the rest, i.e., high babble noise. Some of the background speech came from radio, television, and telephone.

We investigated at spoken-language (Section 3.2) and at speech (Section 3.3). The former was aimed at exploring which spoken-language content changes due to noise contamination, while the latter was aimed at finding the relation between changes in spoken-language levels and those in speech levels. Figure 2 summarizes the levels investigated.

3.2. Analysis at spoken language levels

We examined the impact of noise on the spoken language level from phonetic to syntactic levels.

Phone level: Since the discrete symbols of GSLM can be viewed as pseudo-phones, we defined a counterpart of the WER

Table 2: *UERS [%] of noisy speech*

	SNR			
	15 dB	10 dB	5 dB	0 dB
L-BAB	13.4	17.1	22.0	28.6
M-BAB	20.2	25.9	34.3	47.1
H-BAB	21.8	29.2	40.1	54.1




Figure 3: $P(\text{phoneme}|\text{unit})$. The darker the color, the more uniquely the unit is mapped by a single phoneme.

for the discrete symbol sequences which we call the unit error rate (UER). The UERs were computed in the same manner as the WER but using the discrete symbol sequences obtained from the clean and noisy speech as reference and hypothesis, respectively.

Table 2 shows the average UERs. The UER was above 10 % even under the mildest condition (SNR of 15 dB and L-BAB) and increased significantly as the SNR decreased. This indicates that the GSLM model is vulnerable to noisy inputs when trained with speech without noise. We also observed that background speech deteriorated encoding performance, which was further worsened by lower SNR. This observation suggests that the encoder partially encodes background speech.

Phoneme level: For phoneme-level analysis, we first trained unit-to-phoneme mappings, such as allophone-to-phoneme many-to-one mappings. An allophone is a set of phones mapped to a phoneme in a particular language, e.g., [p^h, p] (phones) → /p/ (phoneme). The IBM model 2 aligner with monotonic alignment constraints aligns unit sequences of clean speech and phoneme sequences obtained from texts in the LibriSpeech corpus. We calculated the probability $P(\text{phoneme}|\text{unit})$ and mapped each unit to its most probable phoneme, e.g., [”128”, ”52”] (units) → /p/ (phoneme). The phonemizer and aligner used were espeak [24] and fast_align library [25], respectively. We calculated the phoneme error rate (PER), just as we calculated the UER in the phone-level evaluation. The reference and hypothesis are phoneme sequences mapped from clean-speech units and noisy-speech units, respectively. Repeated symbols were removed from the sequence to prevent duplication of the same phoneme due to the consecutive estimation of symbols. We used 40 phonemes after removing the stress marks from the espeak notation.

Before discussing the PER, we show alignments from units to phonemes. Figure 3 shows the heatmap of the probability $P(\text{phoneme}|\text{unit})$. We can find three trends. First, some phonemes, e.g., /g, h, j, u, ʃ/, have high probability to their corresponding phonemes, meaning that these phonemes are more uniquely mapped to the single phonemes. However, /n, r, s, ə,

Table 3: PERs [%] of noisy speech

	SNR			
	15 dB	10 dB	5 dB	0 dB
L-BAB	14.8	18.1	21.9	27.2
M-BAB	20.3	25.3	32.5	43.9
H-BAB	24.4	31.9	44.4	61.9




Figure 4: Heatmap of confusion matrix at phoneme level. Values are on \log_{10} scale. Only misrecognized phonemes are counted; correctly recognized ones are not shown in this figure. Noise condition is mildest case: 15 dB SNR of L-BAB noise.

/l/ have many darker cells and mapped from many units. Finally, /q, x, y, ʒ/ have only brighter cells. One individual unit is mapped to each of the phonemes, but the probability is low.

Table 3 lists the average PERs with varying SNR. Consistent with the trend shown in UER, PER increased as SNR decreased and reached closer to H-BAB. In M-BAB and L-BAB, lower PERs (Table 3) were observed compared with UERs (Table 2). However, the PERs in H-BAB increased regardless of the SNR. Therefore, we can deduce the following two points: First, some unit errors in L-BAB and M-BAB are within allophones and negligible at the phoneme level, and second, unit errors in H-BAB change the phonemes. This is reasonable if the background speech in H-BAB causes changes in the phoneme level.

To discuss the PER results, we computed the confusion matrix of the phonemes. Specifically, we calculated the matrix by using `fast_align` to align the groundtruth and hypothesis phonemes used in the PER calculation. The results of the confusion matrix are shown in Figure 4. First, focusing on the groundtruth phonemes (i.e., x-axis), we see that most of the errors are concentrated in specific phonemes, e.g., /e, æ, r/, etc. Therefore, we can say that units corresponding to these phonemes entangle the phonemes. Next, focusing on the error pairs (i.e., the entire matrix), the matrix is very sparse. Some errors are close in articulation, e.g., /r/ (x-axis) \rightarrow /l/ (y-axis); however, most errors are not, e.g., vowels \leftrightarrow consonants. Therefore, it is likely that the GSLM encoder captures features that are different from the human articulations.

Word level: For the word-level analysis, we compared the noisy and resynthesized speech in terms of WER. The references are transcriptions provided by the LibriSpeech corpus,

Table 4: WERs [%] of raw input (before slash) and resynthesized (after slash) speech per SNR. Values of CLEAN are for clean speech

	SNR			
	15 dB	10 dB	5 dB	0 dB
CLEAN	4.3/14.5			
L-BAB	5.4/15.3	5.5/16.5	5.4/20.3	7.9/27.1
M-BAB	5.4/17.1	6.0/20.8	7.7/30.9	12.8/52.5
H-BAB	5.8/19.8	8.5/28.6	11.4/49.0	25.5/75.8




Figure 5: Scatterplots between word frequency and Δ WER per SNR. Blue lines are regression lines.

while the hypotheses are automatic speech recognition (ASR) results obtained from the noisy/resynthesized speech by using Whisper-base [26].

Table 4 lists the average WERs of the noisy and resynthesized speech. Since Whisper had recognition errors on some clean speech, we computed WERs for the clean speech and their resynthesized versions for reference. Comparing the input and resynthesized speech, the difference in WER was above ten points (4.3 \rightarrow 14.5) for the reference. Compared with the reference, the WERs for the noisy inputs increased only a few points in L-BAB and at the SNRs of 10 and 15 dB. The results indicate that GSLM resynthesis dominates in word-level degradation rather than mild non-speech noise contamination.

The WER scores increased significantly under noise conditions with more background speech and lower SNRs, similar to the results in UER and PER. The scores also apparently increased, along with the decrease in SNR and the increase in background speech. These results indicate that noise contamination (particularly, background speech) may change the spoken content in GSLM resynthesis.

For a more detailed analysis, we investigated the relationship between word usage frequency in the corpus and WER. To remove the effects of speech recognition performance, we defined a normalized WER (Δ WER) by subtracting the WERs of the CLEAN noise category from those of L-BAB, M-BAB, and H-BAB for the resynthesized speech.

Figure 5 shows scatterplots between word frequencies and Δ WER per SNRs. The Δ WERs were negatively correlated with word frequency, which became more pronounced as SNR decreased. This result indicates that word substitutions are more likely to occur for less frequent words; for example, words such as “the” appeared the most in the corpus and showed a WER of 3.3 %, but words such as “ill” and “law”, which have the same number of phonemes appeared about 100 times in the entire corpus, and showed a WER of around 40 %. We can also observe that both the slope and bias of the regression line increase as the SNR decreases. This indicates that noise contamination not only degrades the overall performance of word recognition but also worsens the recognition of rare words.

Syntactic level: To quantify syntactic errors by resynthesis, we used a WER counterpart for word class, i.e., part-of-speech,

Table 5: WCERs [%] of raw input (before slash) and resynthesized (after slash) speech per SNR. Values of CLEAN are for clean speech

	SNR			
	15 dB	10 dB	5 dB	0 dB
CLEAN	3.1/11.7			
L-BAB	4.0/12.3	4.0/13.1	4.5/16.1	5.7/21.6
M-BAB	3.9/13.6	4.4/16.5	5.8/24.4	9.9/41.8
H-BAB	4.3/15.6	6.5/23.1	8.8/39.6	20.6/61.8

which we call the word class error rate (WCER). The references and hypotheses are the word class sequences extracted from the texts of the LibriSpeech corpus and the speech recognition results, respectively. The word classes were estimated using the natural language toolkit [27].

Table 5 lists the average WCERs of the noisy and resynthesized speech per SNR. We observe a similar trend in WCER as in UER, PER, and WER. Regardless of the SNRs and noise category, the WCERs were slightly lower than the WERs but still large. This result indicates that the word substitution caused by the resynthesis from the noisy speech frequently changes the word class.

3.3. Analysis at speech level

We next evaluated GSLM at speech levels by quantifying the resynthesis quality of GSLM under noisy environments. For this evaluation, metrics for measuring codec distortion and resynthesized speech quality were used. The former metric is within speech levels, while the latter is a comparison between speech and spoken language levels.

Codec distortion: Since the discrete symbol sequence is a compressed speech representation, GSLM can be viewed as a neural speech codec. Thus, we investigated the distortion between an input speech and its resynthesized version, which we call *codec distortion*. The codec distortion was measured using WARP-Q, which was developed to predict the quality of neural speech codecs [28]. It internally compensates for the temporal discrepancy between the ground truth and resynthesized signal, and is robust to the changes in pronunciation timing after resynthesis that often occur in neural speech codecs. It ranges from one to five points and increases as the codec distortion decreases. We used clean speech as the ground truths for computing WARP-Q.

Figure 6 shows the distributions of the WARP-Q values for noisy and resynthesized speech¹. The WARP-Q improved from the resynthesis as the SNR decreased and the background speech increased which is consistent with the results in Section 3.2. The WARP-Q values of the noisy speech were weakly or not correlated with those of the resynthesized speech under all noise conditions. These results indicate that GSLM resynthesis on noisy speech decrease codec distortion when speech is included in the noise, but does not mitigate distortion in general.

Naturalness: To evaluate the naturalness of the resynthesized speech, we used a pre-trained UTMOS model [29] strong learner [30]. It predicts a five-scale pseudo-mean opinion score

¹We treated the additive noise as a form of codec distortion and computed the WARP-Q values for the noisy speech.




Figure 6: Scatterplots of noisy (raw) and resynthesized (resyn) speech in terms of WARP-Q at various SNRs. Region above diagonal line indicates WARP-Q improvement via GSLM resynthesis.




Figure 7: Scatterplots of Δ WERs and MOS values by UTMOS for resynthesized speech. Blue and orange lines are regression lines for L-BAB and H-BAB, respectively.

(MOS) value from a synthesized speech². Unlike WARP-Q, UTMOS does not require clean speech for computation and is agnostic to spoken-linguistic changes between clean and resynthesized speech. Therefore, we investigated the relation between linguistic changes (Δ WER) and speech quality (UTMOS).

Figure 7 shows the sentence-by-sentence relationship between Δ WER and values by UTMOS. The values by UTMOS were distributed around 3.8 points, and their average decreased as SNR decreased. This result is consistent with the observation in the literature where the discrete speech representation is applied to speech separation/enhancement [32]. There were either weak correlations or none between Δ WER and UTMOS values under all noise conditions, showing that GSLM frequently converts noisy speech into natural but content-altered speech. L-BAB does not degrade the UTMOS value but worsens the WER.

4. Conclusion

We assessed the GSLM performance for noisy speech at the spoken-language and speech levels through speech-resynthesis experiments. The analysis at the spoken-language levels shows that the noise contamination strongly affects resynthesis performance and promotes frequent changes in phones, phonemes, words, and word classes. It also suggests that the GSLM encoder captures different speech features from human articulations. The analysis at the speech levels reveals that GSLM frequently resynthesizes natural but content-altered speech and has the risk of falsification if GSLM is used in real (noisy) situations without any special care.

Acknowledgments. This work is supported by JSPS KAKENHI 22H03639, 23H03418, and Moonshot R&D Grant Number JP-MJPS2011.

²Although the UTMOS was trained with English/Chinese synthesized speech of the Blizzard Challenge or Voice Conversion Challenge, and works robustly for speech synthesized in other languages and with speech synthesis systems [31].

5. References

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proceedings of Advances in Neural Information Processing Systems*, 2017, pp. 6000–6010.
- [2] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional Transformers for language understanding," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, vol. 1, 2019, pp. 4171–4186.
- [3] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 7871–7880.
- [4] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," in *Proceedings of Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 1877–1901.
- [5] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint*, arXiv:1807.03748, 2018.
- [6] Y.-A. Chung, W.-N. Hsu, H. Tang, and J. Glass, "An unsupervised autoregressive model for speech representation learning," in *Proceedings of INTERSPEECH*, 2019, pp. 146–150.
- [7] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, "Mockingjay: Unsupervised speech representation learning with deep bidirectional Transformer encoders," in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2020, pp. 6419–6423.
- [8] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proceedings of Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 12 449–12 460.
- [9] 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.
- [10] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [11] K. Lakhota, E. Kharitonov, W.-N. Hsu, Y. Adi, A. Polyak, B. Bolte, T.-A. Nguyen, J. Copet, A. Baevski, A. Mohamed, and E. Dupoux, "On generative spoken language modeling from raw audio," *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 1336–1354, 2021.
- [12] D. Wells, H. Tang, and K. Richmond, "Phonetic Analysis of Self-supervised Representations of English Speech," in *Proceedings of INTERSPEECH*, 2022, pp. 3583–3587.
- [13] M. de Seyssel, M. Lavechin, Y. Adi, E. Dupoux, and G. Wisniewski, "Probing phoneme, language and speaker information in unsupervised speech representations," in *Proceedings of INTERSPEECH*, 2022, pp. 1402–1406.
- [14] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Sharifi, O. Teboul, D. Grangier, M. Tagliasacchi, and N. Zeghidour, "AudioLM: A language modeling approach to audio generation," *arXiv preprint*, arXiv:2209.03143, 2022.
- [15] 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.
- [16] J. Ao, R. Wang, L. Zhou, C. Wang, S. Ren, Y. Wu, S. Liu, T. Ko, Q. Li, Y. Zhang, Z. Wei, Y. Qian, J. Li, and F. Wei, "SpeechT5: Unified-modal encoder-decoder pre-training for spoken language processing," *arXiv preprint*, arXiv:2110.07205, 2021.
- [17] G. Itai, K. Felix, L. Ann, C. Jade, S. Gabriel, D. Emmanuel, and A. Yossi, "On the robustness of self-supervised representations for spoken language modeling," *arXiv preprint*, arXiv:2209.15483, 2022.
- [18] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Ryan, R. A. Saurous, Y. Ajiomyriannakis, and Y. Wu, "Natural TTS synthesis by conditioning WaveNet on mel-spectrogram predictions," in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2018, pp. 4779–4783.
- [19] W.-N. Hsu, Y.-H. H. Tsai, B. Bolte, R. Salakhutdinov, and A. Mohamed, "Hubert: How much can a bad teacher benefit ASR pre-training?" in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2021, pp. 6533–6537.
- [20] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR corpus based on public domain audio books," in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2015, pp. 5206–5210.
- [21] K. Ito and L. Johnson, "The LJ speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [22] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2019, pp. 48–53.
- [23] J. Thiemann, N. Ito, and E. Vincent, "DEMAND: A collection of multi-channel recordings of acoustic noise in diverse environments," <https://doi.org/10.5281/zenodo.1227121>, 2013.
- [24] "eSpeak NG Text-to-Speech," <https://github.com/espeak-ng/espeak-ng>, 2006.
- [25] C. Dyer, V. Chahuneau, and N. A. Smith, "A simple, fast, and effective reparameterization of IBM model 2," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2013, pp. 644–648.
- [26] A. Radford, K. Jong Wook, X. Tao, B. Greg, M. Christine, and S. Ilya, "Robust speech recognition via large-scale weak supervision," *arXiv preprint*, arXiv:2212.04356, 2022.
- [27] S. Bird, E. Klein, and E. Loper, *Natural language processing with Python: Analyzing text with the natural language toolkit*. O'Reilly Media, Inc., 2009.
- [28] W. A. Jassim, J. Skoglund, M. Chinen, and A. Hines, "Speech quality assessment with WARP-Q: From similarity to subsequence dynamic time warp cost," *IET Signal Processing*, vol. 16, no. 9, pp. 1050–1070, 2022.
- [29] T. Saeki, D. Xin, W. Nakata, T. Koriyama, S. Takamichi, and H. Saruwatari, "UTMOS: UTokyo-SaruLab system for VoiceMOS challenge 2022," in *Proceedings of INTERSPEECH*, 2022, pp. 4521–4525.
- [30] "UTMOS: UTokyo-SaruLab MOS prediction system," <https://github.com/sarulab-speech/UTMOS22>, 2022.
- [31] K. Seki, T. Shinnosuke, S. Takaaki, and H. Saruwatari, "Text-to-speech synthesis from dark data with evaluation-in-the-loop data selection," in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2023, (to appear, arXiv:2210.14850).
- [32] J. Shi, X. Chang, T. Hayashi, Y.-J. Lu, S. Watanabe, and B. Xu, "Discretization and re-synthesis: an alternative method to solve the cocktail party problem," *arXiv preprint*, arXiv:2112.09382, 2021.