

# CoLM-DSR: Leveraging Neural Codec Language Modeling for Multi-Modal Dysarthric Speech Reconstruction

Xueyuan Chen<sup>1</sup>, Dongchao Yang<sup>1</sup>, Dingdong Wang<sup>1</sup>, Xixin Wu<sup>1,3,\*</sup>, Zhiyong Wu<sup>1,2,\*</sup>, Helen Meng<sup>1,2,3</sup>

<sup>1</sup>The Chinese University of Hong Kong, Hong Kong SAR, China

<sup>2</sup>Shenzhen International Graduate School, Tsinghua University, Shenzhen, China

<sup>3</sup>Vocal Engineering Technologies Limited, Hong Kong SAR, China

{xychen, dcyang, ddwang, wuxx, zywu, hmmeng}@se.cuhk.edu.hk

## Abstract

Dysarthric speech reconstruction (DSR) aims to transform dysarthric speech into normal speech. It still suffers from low speaker similarity and poor prosody naturalness. In this paper, we propose a multi-modal DSR model by leveraging neural codec language modeling to improve the reconstruction results, especially for the speaker similarity and prosody naturalness. Our proposed model consists of: (i) a multi-modal content encoder to extract robust phoneme embeddings from dysarthric speech with auxiliary visual inputs; (ii) a speaker codec encoder to extract and normalize the speaker-aware codecs from the dysarthric speech, in order to provide original timbre and normal prosody; (iii) a codec language model based speech decoder to reconstruct the speech based on the extracted phoneme embeddings and normalized codecs. Evaluations on the commonly used UASpeech corpus show that our proposed model can achieve significant improvements in terms of speaker similarity and prosody naturalness<sup>1</sup>.

**Index Terms:** dysarthric speech reconstruction, multi-modal, neural codec, language model

## 1. Introduction

Dysarthria is a prevalent type of speech disorder that is commonly observed in individuals with neuromotor conditions such as Parkinson’s disease and cerebral palsy [1, 2]. This condition results in a significant deterioration in speech quality and voice characteristics from normal speech patterns [3], which greatly hampers dysarthria patients’ daily communication with their family members or caregivers [4]. Dysarthric speech reconstruction (DSR) is a highly effective approach that seeks to improve the speech intelligibility, naturalness, and preserve the original speaker’s timbre by transforming dysarthric speech into normal speech.

The task of DSR is a complex endeavor that has garnered significant research attention. The voice banking-based method collects pre-recorded normal speeches from dysarthric patients before their speech abilities deteriorate to develop personalized text-to-speech (TTS) systems [5], but its applicability is limited to individuals with available normal speech data [6]. The voice conversion (VC) based techniques aim to modify dysarthric speech signals to improve intelligibility and naturalness while preserving the content, such as rule-based VC [7, 8], and statistical VC approaches [9, 10]. Recently, an end-to-end VC based DSR system [11] has been proposed, which involves distilling a speech encoder from a pre-trained automatic speech recognition (ASR) model to replace the text en-

coder in a sequence-to-sequence (seq2seq) TTS system. Compared to a cascaded system that relies on ASR results for TTS, it does not restrict intermediate representations to text characters and can generate speech with lower errors and higher fidelity. Motivated by the prosody and timbre modeling in TTS systems [12, 13, 14], additional components, such as a prosody corrector and speaker encoder, have been introduced to further enhance prosody and speaker similarity [15]. To improve the speech intelligence for patients with severe dysarthria, as well as for speech captured from complex, noisy acoustic environments, a multi-modal framework [16] is first proposed. Two multi-modal encoders are designed and compared to utilize visual information, e.g., lip movements, as additional clues for reconstructing the highly abnormal pronunciations. In order to address the issue of training inefficiency due to complex training strategies and cascaded pipelines, Unit-DSR [17] is proposed to use the discrete speech units extracted from HuBERT [18] for the generation of a normal speech waveform.

Though significant progress has been made, most existing works has focused primarily on improving the speech intelligibility [11, 16, 17]. However, the speaker similarity and prosody naturalness, which are also crucial to a patient’s sense of self-identity and fluent expression, still leave a lot to be desired. In most real-world application scenarios, it is crucial for DSR models to exhibit quick adaptation abilities to new dysarthric patients with limited data, which is difficult for existing speaker encoder based DSR systems [15]. With the development of advanced prompting-based language model (LM) in the field of text analysis [19] and audio processing [20], some zero-shot TTS frameworks have shown strong in-context learning capabilities and diverse outputs with improved speaker similarity and speech naturalness [21, 22, 23], which treat TTS as a language model task with audio codecs as an intermediate representation instead of the traditional mel-spectrogram.

Inspired by the success of neural codec language modeling in zero-shot TTS [21], this paper proposes a codec LM based multi-modal DSR system by leveraging neural codec language modeling with the large, diverse, and multi-speaker normal speech data to improve the reconstruction results, especially for the speaker similarity and prosody naturalness. Firstly, a multi-modal content encoder is adopted to extract the robust phoneme embeddings from dysarthric speech with auxiliary visual inputs. Secondly, we design a speaker codec encoder to extract and modify the speaker-aware codecs from the dysarthric speech, in order to provide the acoustic prompts with original timbre and normal prosody. Thirdly, we use the speech decoder by leverage the neural codec language model to generate the reconstructed speech based on the extracted phoneme embeddings and normal speaker-aware codecs. The contributions of this paper include:

\* Corresponding authors.

<sup>1</sup>Audio samples: <https://Chenxuey20.github.io/CoLM-DSR>




Figure 1: Overall architecture of the proposed CoLM-DSR system with Multi-modal Content Encoder, Speaker Codec Encoder and Codec LM based Speech Decoder.

- We propose the first codec LM based multi-modal DSR system by combining an audio-visual encoder with the neural codec LM framework to reconstruct the dysarthric speech.
- We specially design a novel speaker codec encoder for DSR task by mapping the dysarthric codecs into normal codecs to provide original timbre and normal prosody prompts.
- Both subjective and objective experimental results show that our proposed codec LM based DSR system achieves significant improvements especially for the speaker similarity and prosody naturalness.

## 2. Methodology

Our proposed CoLM-DSR model is illustrated in Figure 1. It mainly consists of a multi-modal content encoder, a speaker codec encoder and a codec LM based speech decoder. Specifically, the multi-modal content encoder strives to extract robust phoneme embeddings from dysarthric audio and visual inputs to provide content prompts. The speaker codec encoder is designed to extract and normalize the speaker-aware codecs from dysarthric speech to provide timbre and prosody prompts. The speech decoder takes phoneme embeddings and speaker-aware codecs as prompt inputs to generate the reconstructed speech.

### 2.1. Multi-modal Encoder for Content Extraction

To reconstruct the linguistic content of original dysarthric speech, a multi-modal encoder is used to extract robust linguistic representations. Following [16], we adopt the multi-modal encoder outputs, i.e., the phoneme probability distribution, as phoneme embeddings, which are denoted as  $\mathbf{p}$ .

As shown in Figure 2 (a), the multi-modal encoder contains an audio feature extractor, a visual feature extractor and an auto-regressive (AR) seq2seq ASR model. Firstly, the log-MMSE speech enhancement algorithm [24] is adopted to reduce the strong background noise, and then 80-dimension filter banks (FBKs)+ $\Delta$  features are extracted as the audio features  $\mathbf{x}^a$ . Secondly, an off-the-shelf face alignment network [25] is employed to detect the lip landmarks, followed by discrete cosine transform (DCT) and linear discriminant analysis (LDA) to downsize and obtain the visual features  $\mathbf{x}^v$ . After that, we further take the common operation of concatenating the temporally aligned audio and visual features along the feature di-

mension, which is followed by a fully-connected (FC) layer to fuse and get the audio-visual features. The audio-visual features are further fed into the following seq2seq ASR model  $\theta_{ASR}$ . It consists of a pretrained AV-HuBERT transformer encoder [26] and an AR decoder with connectionist temporal classification (CTC) [27]. The decoder contains a 512-dimensional location-aware attention module and 2 LSTM layers with 1024 units per layer. Finally, the ASR model outputs the phoneme embeddings  $\mathbf{p}$ , which can be described as

$$\mathbf{p} = f_{ASR}(\mathbf{W}_1(\mathbf{x}^a \oplus \mathbf{x}^v) + \mathbf{z}_1; \theta_{ASR}) \quad (1)$$

where  $\oplus$  is concatenation along the feature dimension,  $\mathbf{W}_1$  and  $\mathbf{z}_1$  are FC-layer parameters.

### 2.2. Speaker Codec Encoder for Timbre Preservation and Prosody Normalization

We specially design a speaker codec encoder by mapping the dysarthric codecs into normal codecs to provide speaker-aware timbre and prosody prompts. As shown in Figure 2 (b), our proposed speaker codec encoder consists of two modules: speaker codec tokenizer and speaker codec normalizer.

**Speaker Codec Tokenizer:** To be specific, we adopt a pretrained neural audio codec model, EnCodec [28], as our tokenizer. EnCodec is a convolutional encoder-decoder model, whose input and output are both 24 kHz audio across variable bitrates. Each embedding produced by the EnCodec encoder is modeled by a residual vector quantization (RVQ) [29], in which the  $\mathbf{8}$  hierarchy quantizers with 1024 entries are finally chosen. Therefore, for each given dysarthric speech  $\hat{\mathbf{s}}$  or normal speech  $\tilde{\mathbf{s}}$ , the corresponding codecs can be obtained and denoted as:

$$\hat{\mathbf{C}}^{T \times 8} = \text{EnCodec}(\hat{\mathbf{s}}), \tilde{\mathbf{C}}^{T \times 8} = \text{EnCodec}(\tilde{\mathbf{s}}), \quad (2)$$

where  $\hat{\mathbf{C}}$  and  $\tilde{\mathbf{C}}$  represent the two-dimensional acoustic code matrix, and  $T$  is the downsampled utterance length.

**Speaker Codec Normalizer:** Since the dysarthric codecs contain abnormal prosodic features and severe noise information, their distribution is different from that of normal codecs, which will seriously affect the prosody and sound quality of reconstructed speech. Therefore, we further design a speaker codec normalizer to map the dysarthric codecs  $\hat{\mathbf{C}}$  into corresponding normal codecs  $\tilde{\mathbf{C}}$  by a speaker verification (SV) estimator  $\theta_{SV}$  to not only preserve the timbre but also modify the prosody. The SV estimator is trained with a generalized end-to-end (GE2E) loss [30], so that the hidden codec representations  $\mathbf{h}$  extracted from codec sequence  $\mathbf{C}$  of the same speaker and different speakers have high and low similarity, respectively. By utilizing a substantial collection of high-quality speech recordings from diverse speakers with varying timbres and natural prosody, we develop a comprehensive normal codec set  $\mathcal{C} = \{\tilde{\mathbf{C}}_i : i = 1, \dots, N\}$ , which can be considered as encompassing the entire space within the domain of normal codecs. Based on this, for any given dysarthric codec sequence  $\hat{\mathbf{C}}$ , we can map it to the closest one  $\tilde{\mathbf{C}}$  in the normal codec set  $\mathcal{C}$  by SV distance, where the SV distance is the L1 distance between the hidden codec representations  $\hat{\mathbf{h}}$  and  $\tilde{\mathbf{h}}_i$  extracted by the SV estimator  $\theta_{SV}$ , formulated as:

$$\hat{\mathbf{h}}_{\sim \hat{\mathbf{C}}} = f_{SV}(\hat{\mathbf{C}}; \theta_{SV}), \tilde{\mathbf{h}}_{i \sim \tilde{\mathbf{C}}_i} = f_{SV}(\tilde{\mathbf{C}}_i; \theta_{SV}) \quad (3)$$

$$\hat{\mathbf{C}} \rightarrow \tilde{\mathbf{C}} = \arg \min_{\tilde{\mathbf{C}}_i \in \mathcal{C}} |\hat{\mathbf{h}}_{\sim \hat{\mathbf{C}}} - \tilde{\mathbf{h}}_{i \sim \tilde{\mathbf{C}}_i}| \quad (4)$$

After that, the corresponding normal speaker-aware codecs  $\tilde{\mathbf{C}}$  with original timbre and natural prosody are obtained.




Figure 2: Diagram of the proposed CoLM-DSR system, where (a), (b) and (c) show the model details of Multi-modal Encoder, Speech Codec Encoder and Neural Codec Language Modeling respectively.

### 2.3. Codec LM based Speech Decoder for Speech Reconstruction

Inspired by the zero-shot TTS [21], we leverage the neural codec LM based decoder to reconstruct the dysarthric speech conditioned on the phoneme embeddings  $\mathbf{p}$  and speaker-aware codecs  $\tilde{\mathbf{C}}$ . The neural codec LM is expected to learn to extract the content and speaker information from the phoneme embeddings and the codecs respectively. Figure 2 (c) shows the process of neural codec language modeling. Specifically, there are two conditional language models in a hierarchical manner.

**AR Transformer Decoder:** During stage 1, an autoregressive (AR) transformer decoder  $\theta_{AR}$  is adopted for the first quantizer  $\tilde{\mathbf{C}}_{:,1}$ , which is conditioned on the phoneme embeddings  $\mathbf{p}$  and the first quantizer of acoustic codecs  $\tilde{\mathbf{C}}_{:,1}$ , formulated as

$$p(\mathbf{C}_{:,1}|\mathbf{p}, \tilde{\mathbf{C}}_{:,1}; \theta_{AR}) = \prod_{t=0}^T p(\mathbf{C}_{t,1}|\mathbf{C}_{<t,1}, \tilde{\mathbf{C}}_{:,1}, \mathbf{p}; \theta_{AR}) \quad (5)$$

**NAR Transformer Decoder:** After obtaining the first quantizer codecs  $\tilde{\mathbf{C}}_{:,1}$  by the AR model and during stage 2-8, a non-autoregressive (NAR) transformer decoder  $\theta_{NAR}$  is used for generating the discrete codecs from the second to the last quantizers,  $\mathbf{C}_{:,j \in [2,8]}$ . It is conditioned on the phoneme embeddings  $\mathbf{p}$ , the speaker-aware codecs  $\tilde{\mathbf{C}}$  and the predicted acoustic tokens belong to the previous codebooks  $\mathbf{C}_{:,<j}$ :

$$p(\mathbf{C}_{:,2:8}|\mathbf{p}, \tilde{\mathbf{C}}; \theta_{NAR}) = \prod_{j=2}^8 p(\mathbf{C}_{:,j}|\mathbf{C}_{:,<j}, \mathbf{p}, \tilde{\mathbf{C}}; \theta_{NAR}) \quad (6)$$

Finally, the whole reconstructed codecs  $\mathbf{C} = \mathbf{C}_{:,1} \oplus \mathbf{C}_{:,2:8}$  with 8 quantizer codes are obtained by the concatenation of each stage result. Then the pre-trained speech codec decoder [28] is used to synthesize the reconstructed speech.

## 3. Experiments

### 3.1. Experimental Settings

Experiments are conducted on the UASpeech [31], VCTK [32] and LibriTTS [33] datasets. The UASpeech corpus is a benchmark disordered speech corpus, which is recorded by an 8-channel microphone array and a video camera with some background noise. We use the VCTK corpus with 105 native speakers to train the SV estimator. The LibriTTS corpus containing 580 hours of normal speech from 2456 speakers is used to

develop the normal codec set and train the Codec LM based speech decoder by teacher-forcing mode. Similar to [16], four speaker-dependent DSR systems are separately built for the four selected speakers (M12, F02, M16 and F04) with the lowest speech intelligibility. Three mel-spectrogram based baseline settings are compared:

- **AON-DSR:** It uses an audio-only encoder to extract phoneme embeddings, a prosody corrector to explicitly model the duration and pitch, a speaker encoder to represent the speaker embedding, and a mel-decoder to reconstruct the mel-spectrogram based on the phoneme and prosody inputs [15].
- **VGG-DSR:** Following AON-DSR system, it uses a VGG-based audio-visual encoder instead of the audio-only encoder to extract phoneme embeddings [16].
- **AVHu-DSR:** Following AON-DSR system, it uses a AVHuBERT-based audio-visual encoder (similar to 2.1) instead of the audio-only encoder to extract phoneme embeddings [16].

All the content encoders are first trained on the whole dysarthric speech datasets of all speakers for 1M steps with batch size of 8, and then finetuned on the target speaker for 2k steps to improve phoneme prediction accuracy. The open-source pre-trained ‘AV-HuBERT Base’ model<sup>1</sup> and ‘EnCodec’ model<sup>2</sup> are adopted in our experiments. The codec LM based decoder is implemented based on an open-source implementation<sup>3</sup> of VALL-E [21], and is trained on 4 NVIDIA V100 GPUs for 300K iterations with a batch size of 4 on each GPU.

### 3.2. Experimental Results

#### 3.2.1. Speaker Similarity Comparison

Subjective tests are conducted to evaluate the speaker similarity of reconstructed speech compared with the original dysarthric speech. 10 subjects are invited to give the 5-point mean opinion score (MOS, 1-bad, 2-poor, 3-fair, 4-good, 5-excellent) for 10 utterances randomly selected from each of four dysarthric speakers, and the scores are averaged and shown in Table 1. As can be observed, all baseline systems still have a poor speaker similarity performance. Compared with the three speaker encoder based baseline systems, our proposed codec LM based

<sup>1</sup>[https://github.com/facebookresearch/av\\_hubert](https://github.com/facebookresearch/av_hubert)

<sup>2</sup><https://github.com/facebookresearch/encodec>

<sup>3</sup><https://github.com/lifeiteng/vall-e>

Table 1: Comparison Results of MOS with 95% Confidence Intervals for Speaker Similarity and Speech Naturalness.

Models	Speaker Similarity				Speech Naturalness			
	M12	F02	M16	F04	M12	F02	M16	F04
Original	-	-	-	-	1.44±0.19	1.70±0.17	2.44±0.22	2.68±0.25
AON-DSR	1.90±0.21	2.61±0.11	2.54±0.18	2.81±0.11	2.90±0.33	2.78±0.31	2.98±0.32	3.12±0.36
VGG-DSR	2.32±0.18	2.78±0.17	2.66±0.21	2.92±0.13	3.00±0.32	2.94±0.29	3.26±0.28	3.20±0.33
AVHu-DSR	2.31±0.21	2.81±0.15	2.69±0.19	3.10±0.20	3.56±0.24	3.62±0.19	3.54±0.19	3.52±0.26
Proposed	<b>3.30±0.17</b>	<b>3.70±0.22</b>	<b>3.58±0.18</b>	<b>3.78±0.19</b>	<b>3.90±0.33</b>	<b>3.80±0.32</b>	<b>3.80±0.33</b>	<b>3.91±0.22</b>

model achieves significant improvements for all the 4 speakers on speaker similarity.

We also employ the speaker verification model [30] as an objective measure to evaluate the speaker similarity between the dysarthric speeches and corresponding reconstructed speeches, and results are shown in Table 2. Our proposed model also achieves the best results for all speakers. Both the subjective and objective results illustrate that our proposed codec prompting based DSR system can preserve more original timbre information benefited from the zero-shot voice clone ability of codec LM. Compared with the speaker encoder based baseline methods requiring large data, our codec LM based model is more suitable for this low-source DSR task.

Table 2: Objective Comparison Results for Speaker Similarity.

Models	M12	F02	M16	F04
AON-DSR	1.135	1.161	1.070	1.061
VGG-DSR	1.137	1.152	1.066	1.069
AVHu-DSR	1.127	1.154	1.074	1.054
Proposed	<b>1.080</b>	<b>1.074</b>	<b>0.987</b>	<b>0.969</b>

### 3.2.2. Speech Naturalness Comparison

To show the speech naturalness improvement of final reconstructed speech compared with the 'Original' dysarthric speech, we also conduct a MOS test on prosody performance. As shown in Table 1, we can see that the original dysarthric speech obtains the lowest score and suffers from very severely abnormal prosody. All DSR systems improve the naturalness of original dysarthric speech, and our proposed CoLM-DSR system achieves the highest scores for all dysarthric patients. The baseline systems rely on predicting the explicit prosodic features (e.g., duration and pitch) from phoneme embeddings during inference, which tends to reconstruct an averaged prosody representation. In contrast, our CoLM-DSR system is trained with large and diverse data in the LM manner, and reconstructed directly based on the acoustic codec prompt during inference, which can generate more natural and diverse prosody effects.

### 3.2.3. Speech Intelligibility Comparison

To show the content intelligibility of final reconstructed speech compared with 'Original' dysarthric speech, we use the publicly released ASR model, Whisper [34], to obtain the word error rate (WER) with greedy decoding. The results are shown in Table 3. Compared with the original dysarthric speech, all the DSR systems achieve a significant improvement. Our proposed CoLM-DSR system also achieves the best results for all speakers. Compared with the SOTA baseline system AVHu-DSR, the WER improvement is not obvious enough, since we adopt the same AVHuBERT-based encoder to extract the content information. It also shows the content encoder is quite important for the speech intelligibility of reconstructed speech.

Table 3: WER Comparison Results for Speech Intelligibility.

Models	M12	F02	M16	F04
Original	98.0%	93.4%	81.6%	67.6%
AON-DSR	65.7%	62.5%	60.3%	55.6%
VGG-DSR	60.6%	59.2%	57.0%	55.4%
AVHu-DSR	55.9%	56.1%	51.3%	48.0%
Proposed	<b>55.6%</b>	<b>55.9%</b>	<b>51.0%</b>	<b>47.9%</b>

### 3.2.4. Investigation on Dysarthric Codecs and Normal Codecs

In order to verify the necessity and effectiveness of our proposed speaker codec normalizer, we further conduct an analysis of the dysarthric codecs and normal codecs. We use the dysarthric codecs as the acoustic prompts directly to generate the reconstructed speech. However, we find that it is difficult to generate normal speech using dysarthric codecs directly for the severe patients, such as M12. Therefore, we only select F04 as an example for comparison. We perform three AB preference tests in terms of audio quality, timbre similarity and prosody naturalness respectively. Listeners are required to select a better utterance for each given utterance pair. The results are shown in Figure 3. We can observe that using normal codecs and using dysarthric codecs achieve relatively consistent performance on timbre similarity. While in terms of audio quality and prosody naturalness, the effect of using normal codecs is significantly better than that of using dysarthric codecs. The results show that abnormal prosodic features and severe noise information seriously affect the prosody naturalness and audio quality of reconstructed speech, and our proposed speaker codec normalizer can effectively not only preserve the original timbre but also improve the prosody and audio quality performance.




Figure 3: Results of AB Preference Test.

## 4. Conclusion

This paper proposes to leverage the neural codec language model for improving dysarthric speech reconstruction results. We combine an audio-visual content encoder with the neural codec language modeling framework. To provide original speaker timbre and natural prosody acoustic prompts, we specially design a normal speaker codec encoder with codec tokenizer and normalizer by mapping the dysarthric codes to normal codecs. Both subjective and objective experimental results on UASpeech corpus show that our proposed CoLM-DSR system can achieve significant improvements especially in terms of speaker similarity and prosody naturalness.

## 5. Acknowledgements

This research is supported by National Natural Science Foundation of China (62076144), Shenzhen Science and Technology Program (WDZC20220816140515001, JCYJ20220818101014 030), CUHK Direct Grant for Research (Ref. No. 4055221), the CUHK Stanley Ho Big Data Decision Analytics Research Centre and the Centre for Perceptual and Interactive Intelligence.

## 6. References

- [1] F. L. Darley, A. E. Aronson, and J. R. Brown, "Motor speech disorders," (*No Title*), 1975.
- [2] T. L. Whitehill and V. Ciocca, "Speech errors in cantonese speaking adults with cerebral palsy," *Clinical linguistics & phonetics*, vol. 14, no. 2, pp. 111–130, 2000.
- [3] F. L. Darley, A. E. Aronson, and J. R. Brown, "Differential diagnostic patterns of dysarthria," *Journal of speech and hearing research*, vol. 12, no. 2, pp. 246–269, 1969.
- [4] A. B. Kain, J.-P. Hosom, X. Niu, J. P. Van Santen, M. Fried-Oken, and J. Staehely, "Improving the intelligibility of dysarthric speech," *Speech communication*, vol. 49, no. 9, pp. 743–759, 2007.
- [5] J. Yamagishi, C. Veaux, S. King, and S. Renals, "Speech synthesis technologies for individuals with vocal disabilities: Voice banking and reconstruction," *Acoustical Science and Technology*, vol. 33, no. 1, pp. 1–5, 2012.
- [6] X. Chen, Q. Huang, X. Wu, Z. Wu, and H. Meng, "Hilvoice: Human-in-the-loop style selection for elder-facing speech synthesis," in *ISCSLP*. IEEE, 2022, pp. 86–90.
- [7] F. Rudzicz, "Adjusting dysarthric speech signals to be more intelligible," *Computer Speech & Language*, vol. 27, no. 6, pp. 1163–1177, 2013.
- [8] S. A. Kumar and C. S. Kumar, "Improving the intelligibility of dysarthric speech towards enhancing the effectiveness of speech therapy," in *2016 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*. IEEE, 2016, pp. 1000–1005.
- [9] S.-W. Fu, P.-C. Li, Y.-H. Lai, C.-C. Yang, L.-C. Hsieh, and Y. Tsao, "Joint dictionary learning-based non-negative matrix factorization for voice conversion to improve speech intelligibility after oral surgery," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 11, pp. 2584–2594, 2016.
- [10] R. Aihara, T. Takiguchi, and Y. Ariki, "Phoneme-discriminative features for dysarthric speech conversion." in *Interspeech*, 2017, pp. 3374–3378.
- [11] D. Wang, J. Yu, X. Wu, S. Liu, L. Sun, X. Liu, and H. Meng, "End-to-end voice conversion via cross-modal knowledge distillation for dysarthric speech reconstruction," in *ICASSP 2020*. IEEE, 2020, pp. 7744–7748.
- [12] X. Chen, C. Song, Y. Zhou, Z. Wu, C. Chen, Z. Wu, and H. Meng, "A character-level span-based model for mandarin prosodic structure prediction," in *ICASSP 2022*. IEEE, 2022, pp. 7602–7606.
- [13] X. Chen, S. Lei, Z. Wu, D. Xu, W. Zhao, and H. Meng, "Un-supervised multi-scale expressive speaking style modeling with hierarchical context information for audiobook speech synthesis," in *COLING*, 2022, pp. 7193–7202.
- [14] X. Chen, X. Wang, S. Zhang, L. He, Z. Wu, X. Wu, and H. Meng, "Stylespeech: Self-supervised style enhancing with vq-vae-based pre-training for expressive audiobook speech synthesis," in *ICASSP 2024*. IEEE, 2024, pp. 12 316–12 320.
- [15] D. Wang, S. Liu, L. Sun, X. Wu, X. Liu, and H. Meng, "Learning explicit prosody models and deep speaker embeddings for atypical voice conversion," *arXiv preprint arXiv:2011.01678*, 2020.
- [16] X. Chen, Y. Wang, X. Wu, D. Wang, Z. Wu, X. Liu, and H. Meng, "Exploiting audio-visual features with pretrained av-hubert for multi-modal dysarthric speech reconstruction," *arXiv preprint arXiv:2401.17796*, 2024.
- [17] Y. Wang, X. Wu, D. Wang, L. Meng, and H. Meng, "Unit-dsr: Dysarthric speech reconstruction system using speech unit normalization," *arXiv preprint arXiv:2401.14664*, 2024.
- [18] 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," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [19] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [20] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Sharifi, D. Roblek, O. Teboul, D. Grangier, M. Tagliasacchi *et al.*, "Audiom: a language modeling approach to audio generation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [21] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li *et al.*, "Neural codec language models are zero-shot text to speech synthesizers," *arXiv preprint arXiv:2301.02111*, 2023.
- [22] Z. Zhang, L. Zhou, C. Wang, S. Chen, Y. Wu, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li *et al.*, "Speak foreign languages with your own voice: Cross-lingual neural codec language modeling," *arXiv preprint arXiv:2303.03926*, 2023.
- [23] S. Ji, J. Zuo, M. Fang, Z. Jiang, F. Chen, X. Duan, B. Huai, and Z. Zhao, "Textrolspeech: A text style control speech corpus with codec language text-to-speech models," *arXiv preprint arXiv:2308.14430*, 2023.
- [24] Y. Ephraim and D. Malah, "Speech enhancement using a minimum mean-square error log-spectral amplitude estimator," *IEEE transactions on acoustics, speech, and signal processing*, vol. 33, no. 2, pp. 443–445, 1985.
- [25] A. Bulat and G. Tzimiropoulos, "How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks)," in *ICCV*, 2017, pp. 1021–1030.
- [26] B. Shi, W.-N. Hsu, K. Lakhotia, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," *arXiv preprint arXiv:2201.02184*, 2022.
- [27] A. H. Liu, H.-y. Lee, and L.-s. Lee, "Adversarial training of end-to-end speech recognition using a criticizing language model," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6176–6180.
- [28] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *arXiv preprint arXiv:2210.13438*, 2022.
- [29] A. Vasuki and P. Vanathi, "A review of vector quantization techniques," *IEEE Potentials*, vol. 25, no. 4, pp. 39–47, 2006.
- [30] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *ICASSP 2018*. IEEE, 2018, pp. 4879–4883.
- [31] H. Kim, M. Hasegawa-Johnson, A. Perlman, J. Gunderson, T. S. Huang, K. Watkin, and S. Frame, "Dysarthric speech database for universal access research," in *Ninth Annual Conference of the International Speech Communication Association*, 2008.
- [32] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, "Superseded-cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2016.
- [33] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "Libritts: A corpus derived from librispeech for text-to-speech," *Interspeech 2019*, 2019.
- [34] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *ICML*. PMLR, 2023, pp. 28 492–28 518.