



Enhancing Dysarthric Speech Recognition for Unseen Speakers via Prototype-Based Adaptation

Shiyao Wang¹, Shiwan Zhao¹, Jiaming Zhou¹, Aobo Kong¹, Yong Qin^{1,*}

¹Nankai University, China

wangshiyao@mail.nankai.edu.cn, zhaosw@gmail.com, zhoujiaming@mail.nankai.edu.cn,
kongaobo9@163.com, qinyong@nankai.edu.cn

Abstract

Dysarthric speech recognition (DSR) presents a formidable challenge due to inherent inter-speaker variability, leading to severe performance degradation when applying DSR models to new dysarthric speakers. Traditional speaker adaptation methodologies typically involve fine-tuning models for each speaker, but this strategy is cost-prohibitive and inconvenient for disabled users, requiring substantial data collection. To address this issue, we introduce a prototype-based approach that markedly improves DSR performance for unseen dysarthric speakers without additional fine-tuning. Our method employs a feature extractor trained with HuBERT to produce per-word prototypes that encapsulate the characteristics of previously unseen speakers. These prototypes serve as the basis for classification. Additionally, we incorporate supervised contrastive learning to refine feature extraction. By enhancing representation quality, we further improve DSR performance, enabling effective personalized DSR.

Index Terms: dysarthric speech recognition, unseen speakers, prototype-based method, supervised contrastive learning

1. Introduction

Dysarthria, a speech disorder caused by various factors such as neuropathy, muscle paralysis affecting speech, decreased muscle contractility, or motor incoordination, is frequently associated with conditions like cerebral palsy, Parkinson’s disease, and head trauma. For those affected, the limited ability to use keyboards or touchscreens makes speech the most convenient means of interacting with devices like smartphones and smart home devices. However, dysarthria-induced changes in breathing, resonance, pronunciation, and prosody significantly impair the performance of speech recognition systems trained on typical speech [1, 2]. Addressing these impairments, dysarthric speech recognition (DSR) technology seeks to bridge the gap for dysarthric speakers, enabling seamless interaction with digital devices. While some studies [3, 4, 5, 6] have focused on developing speaker-independent (SI) DSR models, the inherent variability among dysarthric speakers—due to differences in etiology, age, gender, speaking style, and severity of dysarthria—poses a significant challenge. This variability makes each speaker’s speech patterns distinct, leading to notable performance degradation when SI models are applied to unseen dysarthric speakers.

Consequently, several studies [7, 8, 9, 10, 11] have explored fine-tuning speech recognition models with data from the target speaker for personalized DSR. For instance, Shor et al. [7] achieved this by fine-tuning a conventional speech

recognition model with data from a specific dysarthric speaker, selectively fine-tuning only a subset of the network layers to avoid overfitting. Takashima et al. [8] introduced a two-stage fine-tuning approach, initially leveraging data from multiple dysarthric speakers to develop an SI model that captures general dysarthric speech patterns, followed by further fine-tuning with target speaker data. Recently, Shahamiri et al. [11] devised a novel DSR system that learns to recognize the word shapes spoken by dysarthric speakers and maps them to words. This system achieved optimal performance through the use of target speaker data adaptation. However, these personalized fine-tuning methods require extensive speech data from the target speaker and incur significant training costs. Moreover, the varying severity of dysarthria over time [12] can compromise the long-term efficacy of personalized DSR models, necessitating continuous data collection and model optimization. This not only escalates training expenses but also poses substantial challenges for dysarthric speakers.

To address the aforementioned challenges, we introduce a prototype-based DSR (PB-DSR) method that leverages the unique pronunciation error characteristics of dysarthric speakers. Unlike traditional approaches that necessitate extensive model fine-tuning for each new speaker, our method requires only a minimal dataset to effectively adapt to individual speech patterns. Pronunciation errors, such as phoneme deletion, substitution, insertion, and distortion, are consistent within speakers. Previous studies [13, 14] attempted to address these errors by creating adaptive pronunciation dictionaries for each speaker, based on the analysis of error frequencies. However, these solutions were limited to the most common errors and did not adequately address the misrecognition of severe breath sounds and background noise as phonemes. In contrast, our approach targets word-level pronunciation errors, treating the speech for each word by an individual as a unique class that requires similarity rather than identical matches. Inspired by the principles of prototypical networks for few-shot learning [15], our PB-DSR method utilizes a feature extractor to create per-word prototypes from only few-shot samples, enabling rapid adaptation to the unique speech patterns of each speaker. We utilize the pre-trained HuBERT [16] model, renowned for its general speech recognition capabilities, to extract speech features. To further refine its accuracy for dysarthric speech, HuBERT is initially fine-tuned with a specialized dysarthric speech dataset. Moreover, we integrate supervised contrastive learning (SCL) [17] to enhance feature extraction, adopting a strategy from contrastive learning [18] that optimizes feature representations by minimizing intra-class distances while maximizing inter-class separations, a technique proven effective in various speech tasks [19, 20, 21, 22].

To empirically validate the efficacy of our PB-DSR method,

* Corresponding author.

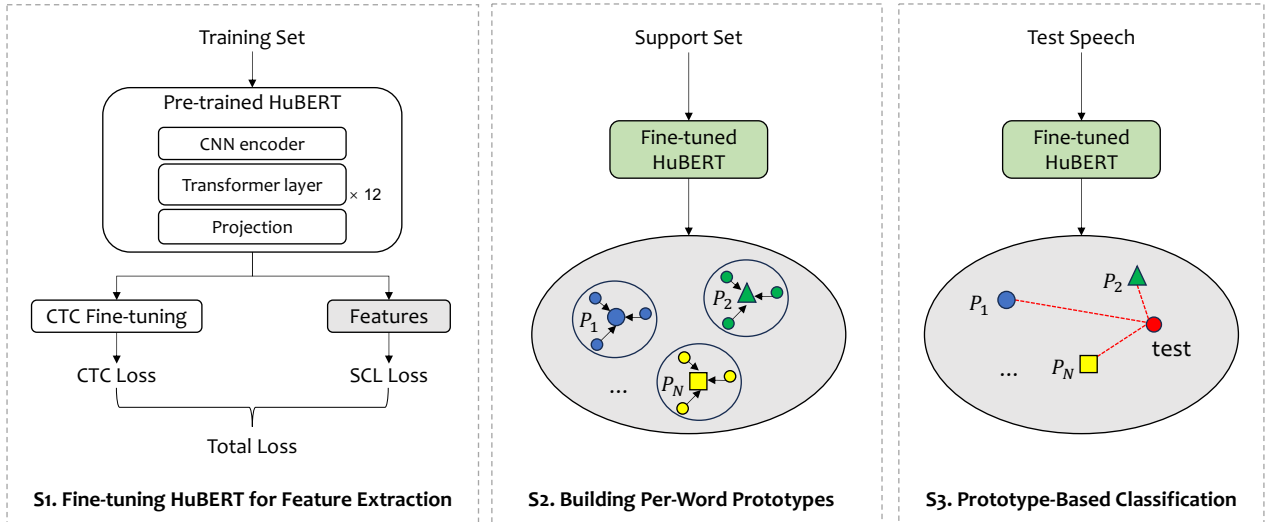


Figure 1: *Prototype-based DSR* comprises three stages: *fine-tuning HuBERT for feature extraction*, *building per-word prototypes*, and *prototype-based classification*.

we conduct experiments with the UASpeech dataset [23]. The results demonstrate notable improvements in speech recognition for unseen dysarthric speakers. Specifically, the PB-DSR method achieves an average absolute reduction in Word Error Rate (WER) of 15.59% compared to its Speaker-Independent (SI) counterpart. Furthermore, incorporating SCL loss into our DSR model training to refine feature extraction leads to an additional 1.21% reduction in WER.

The main contributions of this work are as follows:

- We propose a prototype-based DSR approach that offers a rapid and effective method for improving the recognition of speech from unseen dysarthric speakers.
- We combine CTC loss with SCL loss to train the DSR model to improve performance by learning better feature representations.
- We have demonstrated the effectiveness of the proposed methods on the UASpeech dataset.

2. Proposed Methods

2.1. Prototype-Based DSR

We present an overview of prototype-based DSR (PB-DSR) illustrated in Figure 1, encompassing three key stages: fine-tuning HuBERT for feature extraction, building per-word prototypes, and prototype-based classification.

Stage 1. Fine-tuning HuBERT for Feature Extraction

We fine-tune the pre-trained HuBERT on the DSR task, employing both CTC loss and SCL loss to effectively guide the model’s training during the fine-tuning phase. The fine-tuned HuBERT serves as the feature extractor.

Stage 2. Building Per-Word Prototypes

We classify the speech of different words spoken by the unseen dysarthric speaker as distinct categories, and utilize the limited amount of speech data provided by the unseen dysarthric speaker as a support set. We employ the fine-tuned HuBERT to extract the features of speech from the support set. Subsequently, we average the features of the same word to build per-word prototypes. These prototypes serve as representations of the respective words in the DSR.

Stage 3. Prototype-Based Classification

We employ the fine-tuned HuBERT to extract the feature of test speech. It is important to emphasize that both the test speech and the speech in the support set are from the same unseen dysarthric speaker. Following this, we compute the distance between the test speech feature and each prototype. Ultimately, we choose the word represented by the nearest prototype as the recognition result.

2.2. Combining CTC Loss with SCL Loss

The CTC loss is widely utilized in speech recognition task due to its ability to accommodate data with potential mismatches in alignment between text and speech. Considering the presence of pauses and heavy breathing sounds in dysarthric speech, we opt for the CTC loss function as the primary loss function.

Furthermore, SI models perform poorly in recognizing the speech from speakers with more severe dysarthria, or speech from unseen dysarthric speakers, possibly attributed to the model’s challenge in distinguishing the speech of different words in these speakers. Hence, we integrate the CTC loss with the SCL loss as the total loss function, with the goal of enhancing the model’s ability to learn improved feature representations by simultaneously augmenting inter-class distances and diminishing intra-class distances.

In SCL, within a training batch, samples sharing the identical label are considered positive samples for each other, while those with distinct labels are treated as negative samples. Consequently, the SCL loss, denoted as L_{SCL} , is formulated as follows:

$$L_{SCL} = \sum_{i \in I} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(x_i x_p) / \tau}{\sum_{a \in A(i)} \exp(x_i x_p) / \tau}, \quad (1)$$

where $i \in I = \{1, \dots, N\}$ denotes the index of a speech, $A(i)$ denotes all indices except i , and x_i denotes feature extracted by HuBERT. $P(i)$ denotes all indices of the positive samples of sample i . The features representing the same word as sample i are positive samples. τ is the temperature hyperparameter.

The total loss, denoted as L_{Total} , combines the CTC loss L_{CTC} with the SCL loss L_{SCL} as follows:

$$L_{Total} = L_{CTC} + L_{SCL}. \quad (2)$$

3. Experiments

3.1. Dataset

We assess the efficacy of the proposed method using the UASpeech dataset [23], comprising speech data from 15 dysarthric speakers and 13 control speakers. Details of intelligibility for 15 dysarthric speakers are presented in Table 1, with intelligibility denoting the percentage of speech comprehended by the average listener [24]. The dataset is partitioned into 3 blocks (block 1, 2, and 3), comprising a total of 455 isolated words. Each speaker reads 10 digits, 26 radio alphabet words, 19 computer commands, and 100 common words three times. Each block encompasses one set of these speech. Additionally, there are 300 uncommon words, with each word read once by each speaker. These uncommon words are distributed across the blocks, with each block containing 100 uncommon words. The UASpeech data is recorded with an 8-microphone array, where microphone 1 is utilized for tone synchronization, and microphones 2 to 8 capture the speech signals. Due to data volume requirements for fine-tuning, the process utilizes data from 7 microphones. To expedite PB-DSR and assess its few-shot learning capabilities, per-word prototypes are exclusively built using data solely from microphone 5 (M5).

Below are descriptions of data splits:

- **TRAIN**: All speech from control speakers and speech from block 1 and block 3 of dysarthric speakers.
- **CTEST**: The speech from block 2 of all dysarthric speakers excluding uncommon words speech.

3.2. Experimental Settings

3.2.1. The Settings of DSR Models

We fine-tune HuBERT using the UASpeech dataset under the Fairseq open-source framework¹ to train the DSR model. The DSR model has a parameter size of 94.84M and is trained using a single NVIDIA GeForce RTX 3090. We employ the pre-trained HuBERT Base², which has been pre-trained on the Librispeech 960h dataset³. The DSR model is trained using the configuration file `base_10h.yaml`⁴, employing an initial learning rate of 10^{-5} , 32000 warmup steps, and a batch size of 40. We set the parameter τ in Eq 1 to 0.07. The utilized dictionary comprises the 455 words from the UASpeech dataset, along with the tokens blank, $\langle s \rangle$, $\langle pad \rangle$, $\langle /s \rangle$, and $\langle unk \rangle$. Training will cease if the loss fails to decrease for 10 consecutive epochs. The evaluation metric is the Word Error Rate (WER), measured in percentage.

We first train the vanilla DSR model (denoted as **V**) using the **TRAIN** split, to handle the seen dysarthric speakers. To validate the effectiveness of the proposed PB-DSR method in enhancing the speech recognition performance for unseen dysarthric speakers, we iteratively designate each dysarthric speaker from the UASpeech dataset as ‘unseen’. Subsequently, we remove the data of the designated speaker from the **TRAIN** split, resulting in modified splits denoted as **R-TRAIN**. From this process, we obtain 15 distinct **R-TRAIN**, each correspond-

¹<https://github.com/facebookresearch/fairseq/blob/main/examples/hubert/README.md>

²https://dl.fbaipublicfiles.com/hubert/hubert_base_ls960.pt

³<http://www.openslr.org/12>

⁴https://github.com/facebookresearch/fairseq/blob/main/examples/hubert/config/finetune/base_10h.yaml

Table 1: *Dysarthric speaker intelligibility information*

Intelligibility Level	Speaker ID	Speech Intelligibility (%)
High	M08	93
	M09	86
	M10	93
	F05	95
	M14	90
Mid	M05	58
	F04	62
	M11	62
Low	M07	28
	F02	29
	M16	43
Very low	M01	15
	M04	2
	F03	6
	M12	7

ing to one of the dysarthric speakers in the dataset. We train 15 models on **R-TRAIN** using only the CTC loss function; these models are collectively referred to as **R**. Additionally, we train another set of 15 models employing the proposed combined loss function, which are denoted as **R+**. Subsequently, we apply the prototype-based method to both **R** and **R+** models, resulting in our **PB-DSR** and **PB-DSR+** configurations, respectively.

To further compare our proposed PB-DSR method with the conventional fine-tuning approach, we proceed to fine-tune each **R** model using data from the corresponding unseen speaker. These fine-tuned models are denoted as **FT-R**. Following this, we implement the prototype-based method on **FT-R**, resulting in **FT-R+PB-DSR** configuration. This step allows us to assess the efficacy of the prototype-based method in improving speech recognition performance on fine-tuned models.

3.2.2. The Settings of Prototype-Based DSR

Due to the peak behavior in CTC models [25, 26], coupled with the fact that our DSR model is configured with words as the dictionary, we observe that only the first frame is predicted as a word after the HuBERT feature passes through the CTC layer. Building on this observation, we compare the performance of PB-DSR when building prototypes using the full HuBERT features and solely relying on the first frame of HuBERT features. We find that utilizing only the first frame produces superior outcomes. Therefore, we utilize the first frame of the extracted HuBERT feature to build per-word prototypes.

We utilize the IndexFlatL2 index in the Faiss library [27] to calculate the Euclidean distance between the feature of each test speech and per-word prototypes.

3.3. Results and Discussions

All experimental results are detailed in Table 2. Initially, we benchmark the SI models for seen speakers, trained on the **TRAIN** split, against other studies to validate our model settings and the efficacy of the proposed loss function (Eq. 2). Drawing from an understanding of dysarthric speech features, Bhat et al. [4] devised static and dynamic data augmentation techniques for dysarthric speech, resulting in good speech recognition performance on the UASpeech dataset. As [4] shares our data settings, i.e., evaluating model performance on the **CTEST** split, it serves as a direct comparison. Our models (**V** and **V+**) demonstrates a significant improvement in WER across all intelligibility levels compared to [4], with p-values

Table 2: WER (%) on different settings. “SI” and “SD” denote speaker-independent and speaker-dependent models, respectively. “FT” denotes fine-tuning, and “+” denotes adding SCL loss.

Intelligibility Level	SI for Seen			SI for Unseen		SD w/o FT		SD w/ FT			
	[4]	V	V+	R	R+	PB-DSR	PB-DSR+	[10]	[11]	FT-R	FT-R+PB-DSR
High	6.40	2.35	2.57	10.28	8.27	6.30	5.34	11.80	8.56	4.31	5.12
Mid	14.60	6.01	4.97	26.57	25.27	15.40	14.57	34.67	30.75	5.53	4.89
Low	18.90	7.91	7.23	50.58	45.41	23.21	22.63	26.67	26.75	7.85	6.27
Very Low	61.50	32.08	32.11	84.05	83.22	64.19	61.72	59.50	66.35	39.01	37.67
Average	25.35	12.09	11.72	42.87	40.54	27.28	26.07	33.16	33.10	14.18	13.49

<0.05 , highlighting the enhancements our approach brings to dysarthric speech recognition.

We then compare the DSR performance of **R** and **PB-DSR** for unseen dysarthric speakers. It becomes evident that the **PB-DSR** significantly enhances the DSR performance at each intelligibility level (p-value <0.05). Specifically, it reduces the absolute values of WER by 3.98% (high), 11.17% (mid), 27.37% (low), and 19.86% (very low), respectively. Subsequently, upon comparing **R** with **R+**, it can be observed that the addition of SCL loss enhances the speech recognition performance of DSR model for unseen speakers (p-value <0.05). As the model learns better feature representation, the performance of **PB-DSR** is further enhanced, evidenced by a significant difference (p-value <0.05) between **PB-DSR** and **PB-DSR+**.

We extend our comparison to include recent works by Shahmiri et al. [10, 11], which similarly emphasize target speaker adaptation but with fine-tuning approaches. Notably, even without model fine-tuning, our approach achieves comparable performance to [10, 11] at the very low intelligibility level and outperforms them at other intelligibility levels. It’s important to note, however, that these studies do not employ HuBERT, limiting direct comparisons under identical conditions. Consequently, we compare **PB-DSR** with **FT-R**, where **FT-R** involves fine-tuning. The performance of **PB-DSR** is relatively lower. Nevertheless, this outcome is deemed reasonable given that **PB-DSR** does not require any fine-tuning. Then we compare **FT-R+PB-DSR** with **FT-R**. We observe performance improvements in all intelligibility levels (p-value <0.05) except for the high intelligibility level. This demonstrates that integrating the PB-DSR method with fine-tuned models further enhances DSR performance for unseen dysarthric speakers, underscoring the adaptability and effectiveness of the PB-DSR approach.

3.4. Visualization

We illustrate the correspondence between DSR model performance and speech feature distribution by visualizing the first frame of HuBERT features from speech data. Figure 2 (a) (b) (c) displays the feature distributions for the same speaker under different settings. For a seen speaker, the SI model **V** shows good speech recognition performance, reflected in the clustered pattern of speech features for the same word. Conversely, for unseen speakers, the SI model **R** exhibits poorer performance, evident in a more chaotic feature distribution. However, by integrating SCL, the **R+** model enhances the organization of speech feature distribution, demonstrating its potential to improve model performance across different speaker scenarios.

In Figure 2 (d), we examine the speech recognition outcomes for an unseen speaker using **R+** and **PB-DSR+**. It’s evident that features misrecognized by the **R+** model are nearly

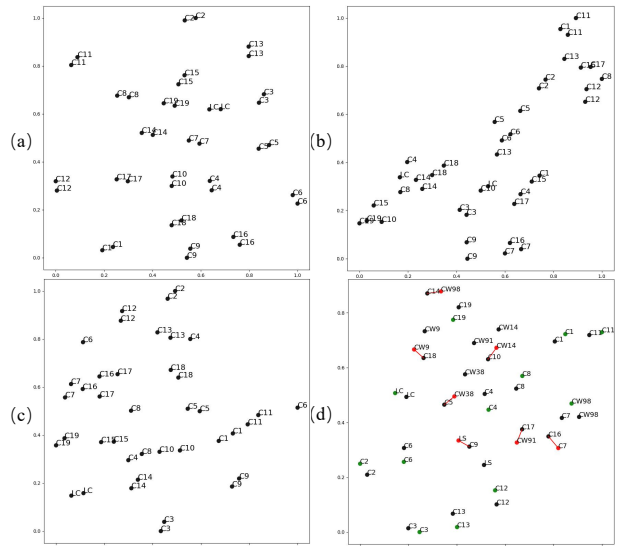


Figure 2: Visualizations of speech feature distributions: (a) Seen speaker in **V**. (b) Unseen speaker in **R**. (c) Unseen speaker in **R+**. (d) Enhancing **R+** by **PB-DSR+**. Black dots in (d) represent prototypes, while green and red dots denote samples correctly and incorrectly classified by **R+**, respectively, with their labels from **R+** predictions. A red line connects a red dot to its correct prototype by **PB-DSR+**. Each label is the word ID in the UASpeech dataset.

aligned with their correct prototypes, showcasing proximity but a crucial lack of precise recognition due to the model’s inability to adjust to the novel speech patterns of the unseen speaker. Conversely, **PB-DSR+**, by employing straightforward distance metrics for classification, accurately recognizes these samples. This methodological refinement significantly boosts speech recognition accuracy for unseen speakers, highlighting the efficacy of PB-DSR in overcoming the adaptation challenges faced by conventional SI models.

4. Conclusions

We propose a prototype-based DSR method that effectively enhances the speech recognition performance for unseen dysarthric speakers without the need for fine-tuning or substantial data from those speakers. Additionally, we introduce supervised contrastive learning to enhance the performance of DSR. Comprehensive experiments on the UASpeech dataset affirm the effectiveness of our proposed method.

5. Acknowledgments

This work has been supported in part by NSF China (Grant No.62271270).

6. References

- [1] M. Tu, A. Wisler, V. Berisha, and J. M. Liss, "The relationship between perceptual disturbances in dysarthric speech and automatic speech recognition performance," *The Journal of the Acoustical Society of America*, vol. 140, no. 5, pp. EL416–EL422, 2016.
- [2] L. De Russis and F. Corno, "On the impact of dysarthric speech on contemporary asr cloud platforms," *Journal of Reliable Intelligent Environments*, vol. 5, pp. 163–172, 2019.
- [3] A. Hernandez, P. A. Pérez-Toro, E. Nöth, J. R. Orozco-Arroyave, A. K. Maier, and S. H. Yang, "Cross-lingual self-supervised speech representations for improved dysarthric speech recognition," in *Interspeech*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:24794022>
- [4] C. Bhat, A. Panda, and H. Strik, "Improved asr performance for dysarthric speech using two-stage data augmentation," *Proc. Interspeech 2022*, pp. 46–50, 2022.
- [5] T. Wang, S. Hu, J. Deng, Z. Jin, M. Geng, Y. Wang, H. Meng, and X. Liu, "Hyper-parameter Adaptation of Conformer ASR Systems for Elderly and Dysarthric Speech Recognition," in *Proc. INTERSPEECH 2023*, 2023, pp. 1733–1737.
- [6] E. Kim, Y. Chae, J. Sim, and K. Lee, "Debiased Automatic Speech Recognition for Dysarthric Speech via Sample Reweighting with Sample Affinity Test," in *Proc. INTERSPEECH 2023*, 2023, pp. 1508–1512.
- [7] J. Shor, D. Emanuel, O. Lang, O. Tuval, M. Brenner, J. Cattiau, F. Vieira, M. McNally, T. Charbonneau, M. Nollstadt *et al.*, "Personalizing asr for dysarthric and accented speech with limited data," *arXiv preprint arXiv:1907.13511*, 2019.
- [8] R. Takashima, T. Takiguchi, and Y. Arikawa, "Two-step acoustic model adaptation for dysarthric speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6104–6108.
- [9] D. Wang, J. Yu, X. Wu, L. Sun, X. Liu, and H. Meng, "Improved end-to-end dysarthric speech recognition via meta-learning based model re-initialization," in *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2021, pp. 1–5.
- [10] S. R. Shahamiri, V. Lal, and D. Shah, "Dysarthric speech transformer: A sequence-to-sequence dysarthric speech recognition system," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2023.
- [11] S. R. Shahamiri, K. Mandal, and S. Sarkar, "Dysarthric speech recognition using depthwise separable convolutions: Preliminary study," in *2023 International Conference on Speech Technology and Human-Computer Dialogue (SpeD)*. IEEE, 2023, pp. 78–82.
- [12] K. Tomanek, K. Seaver, P.-P. Jiang, R. Cave, L. Harrell, and J. R. Green, "An analysis of degenerating speech due to progressive dysarthria on asr performance," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [13] K. T. Mengistu and F. Rudzicz, "Adapting acoustic and lexical models to dysarthric speech," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2011, pp. 4924–4927.
- [14] Y. Sawa, R. Takashima, and T. Takiguchi, "Adaptation of a pronunciation dictionary for dysarthric speech recognition," in *2022 IEEE 4th Global Conference on Life Sciences and Technologies (LifeTech)*. IEEE, 2022, pp. 631–635.
- [15] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," *Advances in neural information processing systems*, vol. 30, 2017.
- [16] 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.
- [17] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," *Advances in neural information processing systems*, vol. 33, pp. 18 661–18 673, 2020.
- [18] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [19] L. Wu, D. Zong, S. Sun, and J. Zhao, "A sequential contrastive learning framework for robust dysarthric speech recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7303–7307.
- [20] X. Wang, S. Zhao, and Y. Qin, "Supervised Contrastive Learning with Nearest Neighbor Search for Speech Emotion Recognition," in *Proc. INTERSPEECH 2023*, 2023, pp. 1913–1917.
- [21] J. Zhou, S. Zhao, N. Jiang, G. Zhao, and Y. Qin, "Madi: Inter-domain matching and intra-domain discrimination for cross-domain speech recognition," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [22] H. Wang, X. Zheng, and Y. Qin, "Intermediate-task learning with pretrained model for synthesized speech mos prediction," in *2023 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2023, pp. 378–383.
- [23] 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.
- [24] H. Chandrashekar, V. Karjigi, and N. Sreedevi, "Investigation of different time-frequency representations for intelligibility assessment of dysarthric speech," *Ieee transactions on neural systems and rehabilitation engineering*, vol. 28, no. 12, pp. 2880–2889, 2020.
- [25] A. Zeyer, R. Schlüter, and H. Ney, "Why does ctc result in peaky behavior?" *arXiv preprint arXiv:2105.14849*, 2021.
- [26] J. Zhou, S. Zhao, Y. Liu, W. Zeng, Y. Chen, and Y. Qin, "Knn-ctc: Enhancing asr via retrieval of ctc pseudo labels," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 11 006–11 010.
- [27] J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with gpus," *IEEE Transactions on Big Data*, vol. 7, no. 3, pp. 535–547, 2019.