



Debiased Automatic Speech Recognition for Dysarthric Speech via Sample Reweighting with Sample Affinity Test

Eungbeom Kim^{1*}, Yunkee Chae^{1*}, Jaeheon Sim¹, Kyogu Lee^{1,2}

¹IPAI, ²AIIS, Seoul National University, Seoul, Republic of Korea

{eb.kim, yunkimo95, sjhoney0112, kglee}@snu.ac.kr

Abstract

Automatic speech recognition systems based on deep learning are mainly trained under empirical risk minimization (ERM). Since ERM utilizes the averaged performance on the data samples regardless of a group such as healthy or dysarthric speakers, ASR systems are unaware of the performance disparities across the groups. This results in biased ASR systems whose performance differences among groups are severe. In this study, we aim to improve the ASR system in terms of group robustness for dysarthric speakers. To achieve our goal, we present a novel approach, sample reweighting with sample affinity test (Re-SAT). Re-SAT systematically measures the debiasing helpfulness of the given data sample and then mitigates the bias by debiasing helpfulness-based sample reweighting. Experimental results demonstrate that Re-SAT contributes to improved ASR performance on dysarthric speech without performance degradation on healthy speech.

Index Terms: speech recognition, debiasing, dysarthric speech

1. Introduction

Automatic speech recognition (ASR) performance across groups should be fair regardless of disorder, race, age, and dialect for a trustworthy and inclusive system. Although ASR has been improved along with the success of deep learning, ASR systems based on deep learning tend to be vulnerable to biases [1, 2]. Consequently, the bias in ASR systems interferes with the ASR system’s trustworthiness by causing poor worst group performance. Applying ASR on dysarthria, which is a type of motor speech disorder, also suffers from the lack of group robustness between healthy and dysarthric speakers. In other words, a dysarthric speaker group has a lower performance than a healthy speaker group on ASR. This is because dysarthric speech has low intelligibility, whereas deep learning has a tendency to be easily fitted to shortcuts [3, 4]. However, most of the previous dysarthric speech recognition studies mainly focus on the sole performance of the ASR system on dysarthric speech. In this study, we aim to improve the performance of ASR on dysarthric speech from a debiasing point of view. To the best of our knowledge, this is the first work to endeavor to make a debiased ASR system for dysarthric speech.

In this regime, we focus on a debiasing method based on sample reweighting. Recently, sample reweighting such as [5, 6] is proposed as a promising solution to handle the bias

problem. Sample reweighting methods mainly consist of 1) estimating the helpfulness of a sample for debiasing and 2) reweighting the sample based on helpfulness. To estimate the debiasing helpfulness, [5] defines *bias-conflicting sample* and *bias-aligned sample* which denote an incorrectly classified sample and a correctly classified sample from an unintended decision rule of a biased model, respectively. The strategies utilizing those definitions are based on the assumption that upweighting the data samples with a large loss from the biased model, i.e. bias-conflicting samples, leads to improved generalization performance of underfitted groups. This assumption, however, hardly holds in real-world problems. For example, although an outlier might have a large loss value, it is an undesirable strategy to upweight the outlier.

Therefore, we define a new taxonomy to directly categorize the samples by debiasing helpfulness; *bias-blocking sample* denotes a sample that mitigates the model’s bias, and *bias-accelerating sample* denotes a sample that accelerates the model’s bias when the model is trained on each sample. To estimate whether a given sample is *bias-blocking sample*, we propose a sample affinity test (SAT) which systematically measures debiasing helpfulness. SAT is based on the novel metric dubbed sample affinity, which denotes a training effect of the given sample to the other samples’ loss. This is inspired by task affinity [7] which estimates inter-task affinity using the loss shift of each task after lookahead training of a given task. Intuitively, SAT measures debiasing helpfulness using sample affinity of the given sample on the bias-conflicting sample set, based on a one-step lookahead training of a given sample. Unlike the loss-based methods, SAT can filter unhelpful bias-accelerating samples even if the samples are bias-conflicting.

Based on SAT, we propose a novel sample reweighting method, Re-SAT, to challenge the debiasing problem of dysarthric speech recognition. Re-SAT consists of four sequential components for implementation. The first component estimates bias-conflicting samples based on loss, following [6]. Secondly, the SAT component is activated based on the estimated bias-conflicting samples to accurately measure the debiasing effect of each sample. In the third component, the SAT result is normalized through sorting and mapped to the weights. Finally, Re-SAT trains the models with the reweighted samples.

In summary, we observe the biased performance of ASR on disordered speech and analyze the ASR system through the lens of debiasing. Furthermore, we present sample affinity test (SAT) to directly measure the debiasing helpfulness of samples. We also propose Re-SAT, a novel method for sample reweighting that is available under the group label-free environment and mitigates the margin between the ASR performances of healthy speakers and dysarthric speakers.

* Equal contribution. This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) [No. 2022-0-00641, XVoice: Multi-Modal Voice Meta Learning, 90%] and [NO.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University), 10%].

2. Related work

Dysarthric speech recognition ASR on dysarthric speech still remains a challenging problem. Since the disordered speech dataset has limited scalability due to the difficulty of data collection, data augmentation on dysarthric speech has been widely studied [8, 9, 10, 11, 12]. Other investigations adopt pre-trained self-supervised learning models [13, 14, 15, 16] for improved disordered speech recognition. On the other hand, [17, 18, 19] utilize some prior knowledge that a dysarthric speaker has distinct characteristics compared to healthy or other dysarthric speakers by leveraging speaker adaptation using the spectro-temporal level. In this study, we approach ASR on dysarthric speech from a different point of view, debiasing, which has never been used to the best of our knowledge.

Debiasing Debiasing is a challenging but essential area for fair and inclusive deep learning. Group robustness, which denotes an ability to perform satisfactorily across different groups in terms of the given task, is one of the key factors of a debiased model. We focus on group robustness without group annotations, for a wider real-world application. Sample reweighting is one of the promising methods for debiasing. While empirical risk minimization (ERM) uniformly averages all of the performances across the given samples, sample reweighting methods such as Learning from Failure (LfF) [5] and Just Train Twice (JTT) [6] upweight the samples that are expected to belong to the poor performance group called bias-conflicting samples. These methods are based on the intuition that upweighting the bias-conflicting samples with large losses leads to a debiased model, which is not always true. In particular, we test ASR systems on dysarthric speakers, each of who can be regarded as a respective group due to the variety of dysarthria. In this complex real-world environment, reweighting the samples using a fine-grained metric is an important issue. For this reason, we aim to propose direct and accurate criteria for reweighting beyond the estimation of bias-conflicting samples based on loss.

Task affinity Given a set of tasks in a multi-task learning setup, task grouping aims to find the clusters of tasks that should be trained simultaneously for improved performance. Task affinity [7] is proposed to address the task grouping problems. To examine the task affinity of task i on task j , they leverage a lookahead update on task i . Then, they compute the loss shift after the lookahead update to explore the effect of task i on j . Inspired by task affinity, we propose sample affinity to investigate the inter-sample effect for debiasing.

3. Method

In this section, we introduce our sample reweighting with sample affinity test (Re-SAT), which is proposed to challenge the ASR on dysarthric speech. Re-SAT consists of four components: 1) bias-conflicting sample estimation, 2) sample affinity test, 3) normalized reweighting and 4) training, as shown in Figure 1. The data samples are reweighted through the first three components, and then Re-SAT trains the model with the reweighted data samples in the 4) training block. In summary, Re-SAT upweights the bias-blocking samples and downweights the bias-accelerating samples using sample affinity test for debiasing. Details are as follows.

3.1. Bias-conflicting sample estimation

At the first step, Re-SAT estimates the bias-conflicting samples using the losses of samples in each batch. Re-SAT regards the samples with the largest K losses in a batch as the bias-

conflicting samples. Unlike the previous research [6] which fixes the estimation results, Re-SAT keeps estimating the bias-conflicting samples during training to accurately reflect the current state of the model. That is, the estimation result for each sample can be modified throughout training.

Although Re-SAT estimates the bias-conflicting samples, Re-SAT does not map the estimation results for reweighting because the large loss does not guarantee the debiasing ability, as we introduced in the above sections. We address this issue by designing a novel criterion that directly estimates the bias-blocking samples, for debiasing ability.

3.2. Sample affinity test

We propose a novel sample affinity test (SAT) to evaluate the debiasing ability of samples to filter the bias-accelerating samples. Given the model f_θ , the samples in the mini-batch x_1, \dots, x_N , and the estimated bias-conflicting samples $\hat{b}_1, \dots, \hat{b}_K$, SAT utilizes single step lookahead updating with respect to x_i as

$$\theta_{x_i}^{LA} = \theta - \eta \nabla L(x_i; f_\theta), \quad (1)$$

for $i = 1, \dots, N$ where η is learning rate, N is batch size, and L is a loss function. After the lookahead updating process, SAT compares the loss of 1) the lookahead model $\theta_{x_i}^{LA}$ and 2) the original model θ on the bias-conflicting samples to calculate the averaged sample affinity:

$$SA(x_i) = \frac{1}{K} \sum_{\forall k} \left(1 - \frac{L(\hat{b}_k; f_{\theta_{x_i}^{LA}})}{L(\hat{b}_k; f_\theta)} \right). \quad (2)$$

We present sample affinity inspired by task affinity [7] for debiasing. Note that SAT computes the sample affinity every step to take into account the current state of the model, while task affinity is averaged across the whole training step.

3.3. Normalized reweighting

Sample affinity on the bias-conflicting samples approximates the bias-blocking effect of a given sample. However, sample affinity is not an absolute but relative score which depends on the bias-conflicting sample set, the current state of the model, and the maturity of the model. Therefore, Re-SAT does not directly map the sample affinity to weights. Instead, Re-SAT normalizes the sample affinity by extracting the rank of the sample affinity with respect to the descending order within the batch. This normalization stabilizes unwanted shifts in the learning rate.

For sorted samples $x_1, \dots, x_r, \dots, x_N$ with respect to sample affinity in descending order, Re-SAT finally reweights the samples using the function $w : \mathbb{N} \rightarrow \mathbb{R}^+$, which is defined as follows:

$$w(r) = \frac{\exp(s(N-r)/(N-1))}{\sum_{r=1}^N \exp(s(N-r)/(N-1))} \quad (3)$$

for $r = 1, \dots, N$ where r is the rank of the sample x_r and s is the constant hyperparameter. For our experiments, we set batch size $N = 32$ and $s = 4$.

3.4. Training

In this component, Re-SAT trains the model f_θ with the reweighted data samples. Unlike the vanilla training that averages the loss of each data sample, Re-SAT leverages weighted

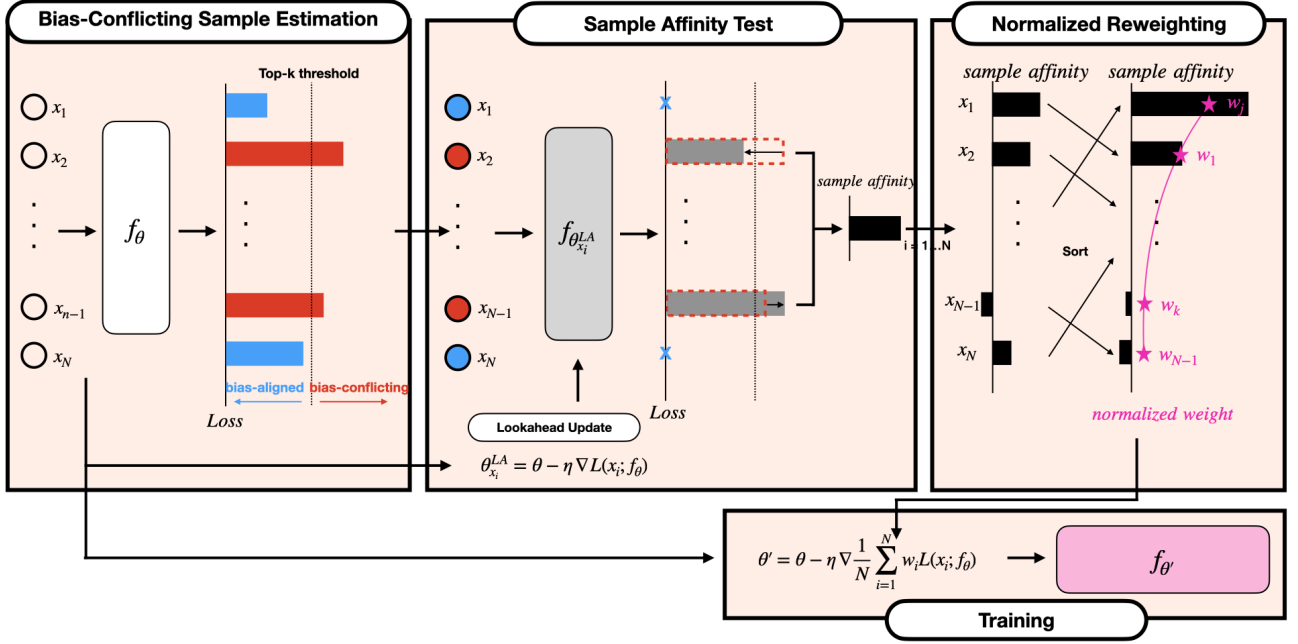


Figure 1: Illustration of Re-SAT for debiased dysarthric speech recognition. Re-SAT consists of 1) bias-conflicting sample estimation, 2) sample affinity test, 3) normalized reweighting, and 4) training blocks.

average for total loss as:

$$\theta' = \theta - \eta \nabla \frac{1}{N} \sum_{r=1}^N w(r) L(x_r; f_\theta). \quad (4)$$

for sorted samples x_1, \dots, x_N with respect to sample affinity $SA(x_i)$ in descending order.

4. Experiments

4.1. Dataset

UASpeech corpus [20] is used for our experiments. UASpeech corpus is one of the largest English dysarthric speech datasets, composed of speech from 15 dysarthric speakers and 13 healthy control speakers. Each dysarthric speaker is classified into very low, low, mid, and high levels of intelligibility. The dataset consists of three blocks B1, B2, and B3. Each block contains 155 common words and 100 uncommon words spoken by dysarthric and healthy control speakers. We used block 1 and block 3 of all healthy and dysarthric speech recorded on microphone M5 as a training set, and all data in block 2 as the test set. Note that we used not only dysarthric speech but also healthy speech as the test set. Basically, UASpeech provides the speech data denoised by *noisereduce* [21, 22]; therefore we used this version for our experiments. We excluded the uncommon words from both the training and test set and no additional data augmentation was conducted.

4.2. Model

We fine-tuned *Whisper* [23], the recent state-of-the-art ASR model, trained on a large amount of labeled audio-transcription data. It employs a simple encoder-decoder Transformer architecture. We used the *Whisper-tiny*, pre-trained on English only with 39M parameters, which is available at HuggingFace transformers repository [24].

4.3. Experimental Setups

We trained the model using AdamW optimizer [25] with a learning rate of 10^{-5} , weight decay of 0.1, and batch size of 32 for 30 epochs. For the proposed method, Re-SAT, we investigated the effect of the number of bias-conflicting samples in batch as $K \in \{2, 4, 8, 16\}$. For comparison, we tested the JTT method [6] with upweighting value 25 and bias-conflicting sample identification epoch 3. For the ablation studies, we tested the model named Re-Loss, which utilizes the loss-based sample reweighting instead of sample affinity and follows Re-SAT in other settings, to figure out the impact of the sample affinity test.

5. Results

We investigated with the effect of the proposed method, Re-SAT, by comparing it to empirical risk minimization (ERM) and JTT for debiasing. In contrast to most of the dysarthric speech recognition studies that focus only on dysarthric speech, we investigated the effect of Re-SAT to ASR system on healthy speech for a fair comparison. This is important for debiasing because the performance disparity on dysarthric speech and healthy speech determines the group robustness of the systems.

Results in Table 1 show word error rate (WER) of the models on each speaker and intelligibility group. It is shown that Re-SAT contributes to performance improvements for each group compared to ERM by relatively decreasing 12.02% (54.66 \rightarrow 48.09), 11.10% (19.19 \rightarrow 17.06), 7.22% (13.57 \rightarrow 12.59), and 13.06% (6.43 \rightarrow 5.59) in terms of the averaged WER for the very low, low, mid, and high intelligibility group, respectively. Surprisingly, Re-SAT also surpasses ERM on the healthy speaker group. From a speaker-wise view, it is observed that the performances of 12 out of 15 speakers show enhanced results in terms of WER on the Re-SAT-based ASR model. These results demonstrate the robust performance gain of Re-SAT even un-

		WER (%)			
	Speaker (Intelligibility %)	ERM	JTT	Re-Loss	Re-SAT (Ours)
VL	M04 (2%)	84.39	79.87	82.45	80.26
	F03 (6%)	44.33	42.21	40.28	38.25
	M12 (7%)	44.41	42.04	41.72	35.38
	M01 (17%)	50.97	51.77	47.10	44.19
	Avg.	54.66	52.46	51.50	48.09
L	M07 (28%)	16.22	23.78	16.96	16.04
	F02 (29%)	19.72	22.86	16.59	15.94
	M16 (43%)	22.04	23.55	20.32	19.57
	Avg.	19.19	23.39	17.84	17.06
M	M05 (58%)	12.63	13.55	12.90	12.72
	M11 (62%)	14.30	12.80	12.04	10.11
	F04 (62%)	13.88	13.88	13.22	14.64
	Avg.	13.57	13.44	12.75	12.59
H	M09 (86%)	9.86	11.15	8.66	8.76
	M14 (90%)	12.17	13.73	11.98	9.86
	M10 (93%)	3.50	3.32	2.40	2.76
	M08 (95%)	4.61	5.53	4.70	3.78
	F05 (95%)	2.03	4.33	2.03	2.76
Avg.	6.43	7.61	5.95	5.59	
Avg. (Dysarthric)		21.49	22.25	20.15	19.93
Avg. (Healthy)		3.81	4.21	3.75	3.40
Avg. (Overall)		12.93	13.51	12.20	12.08

Table 1: ASR results in terms of Word Error Rate (WER) on UASpeech corpus. VL/L/M/H refer to very low/low/mid/high intelligibility groups, respectively. The percentage of intelligibility is measured based on how accurately naive human listeners can transcribe isolated words produced by speakers, according to [20]. K of Re-SAT is set to 4.

der a diverse group environment. On the contrary, the other debiasing method, JTT, increases the averaged WER on healthy speech, low and high intelligibility speech although JTT reduces the very low and mid intelligibility group’s WER, compared to those of ERM.

Figure 2 shows the rank sorted in descending order of training loss within the batch with respect to each intelligibility group. For very low, low, and mid-intelligibility groups that are vulnerable to bias, the increasing and large rank value is desired because it means that the loss for the biased group is decreased, even though it is approximated by training loss. Interestingly, debiasing of JTT works only on the single worst (very low) intelligibility group and even deteriorates the low and mid groups’ ranks. This problem can be caused by the structure of JTT’s 1) binary classification of bias-conflicting samples and 2) which are fixed at an early stage. On the other hand, Re-SAT leverages 1) fine-grained reweighting and 2) updated bias-conflicting set through training and shows desirable results across diverse groups. Figure 2 also demonstrates that ERM is inappropriate for debiasing. Although the ranks on very low, low, and mid intelligibility groups of ERM increase at the early stage for all training schemes, the ranks decrease at the late stage of ERM. Even worse, the healthy speaker’s rank shows the highest rank compared to other methods, which means the biased performance of ERM. We also explore the Re-Loss, which substitutes the sample affinity-based reweighting of

		Intelligibility group					
K		VL	L	M	H	Avg. (D)	Avg. (H)
ERM		54.66	19.19	13.57	6.43	21.49	3.81
2		50.12	17.49	14.15	5.75	19.93	3.72
4		48.09	17.06	12.59	5.59	19.05	3.40
8		49.44	17.61	12.62	6.36	19.75	3.84
16		55.31	21.10	13.60	6.91	22.21	3.96
	Avg.	12.93	13.51	12.20	5.59	19.93	3.40

Table 2: Comparison on hyperparameter K with word error rate (WER). D and H refer to dysarthric and healthy speech, respectively.




Figure 2: Illustrations of the averaged loss ranks sorted by descending order in the mini-batch for each intelligibility group. The x-axis denotes training epochs and the y-axis denotes the loss rank in mini-batch.

Re-SAT with loss-based reweighting, to investigate the importance of the sample affinity test. Although Re-Loss outperforms ERM and JTT by a large margin, Re-SAT still shows a surpassing performance. This demonstrates the precise bias-blocking sample estimation ability of Re-SAT, as we expected.

Intuitively, selecting decent bias-conflicting samples takes an important role in Re-SAT. For this reason, we investigate the size of the selected bias-conflicting sample set K in Re-SAT as shown in Table 2. Re-SAT shows stable improvements across the intelligibility groups where $K = 2, 4, 8$ compared to ERM for all intelligibility groups. This supports our design choice of Re-SAT which re-estimates the bias-conflicting samples through training. Even though Re-SAT selects small bias-conflicting samples, Re-SAT can re-estimate the new underfitted data sample as a bias-conflicting sample through training. On the contrary, Re-SAT with $K = 16$ even degrades the ASR results, since setting a large K as 16 interferes with the accurate estimation of the bias-conflicting samples.

6. Conclusions

In this study, we address a debiasing problem of ASR on dysarthric speech towards a fair and inclusive ASR system. In contrast to the previous research that focuses only on the ASR performance for dysarthric speakers, we explore the fairer validation system by analyzing the performance of healthy and dysarthric speech at the same time. To achieve our goal, we propose a novel debiasing method based on sample reweighting, Re-SAT. The ASR system using Re-SAT surpasses the other baselines across the diverse dysarthric speakers and shows a robust performance gain over various hyperparameters. For future work, it is worth exploring the integrated system of Re-SAT and other dysarthric speech recognition systems by leveraging the versatile structure of Re-SAT.

7. References

- [1] J. L. Martin and K. E. Wright, "Bias in automatic speech recognition: The case of african american language," *Applied Linguistics*, 2022.
- [2] S. Feng, O. Kudina, B. M. Halpern, and O. Scharenborg, "Quantifying bias in automatic speech recognition," *arXiv preprint arXiv:2103.15122*, 2021.
- [3] S. Beery, G. Van Horn, and P. Perona, "Recognition in terra incognita," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 456–473.
- [4] S. Asgari, A. Khani, F. Khani, A. Gholami, L. Tran, A. Mahdavi-Amiri, and G. Hamarneh, "Masktune: Mitigating spurious correlations by forcing to explore," in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.
- [5] J. Nam, H. Cha, S. Ahn, J. Lee, and J. Shin, "Learning from failure: De-biasing classifier from biased classifier," *Advances in Neural Information Processing Systems*, vol. 33, pp. 20 673–20 684, 2020.
- [6] E. Z. Liu, B. Haghgoo, A. S. Chen, A. Raghunathan, P. W. Koh, S. Sagawa, P. Liang, and C. Finn, "Just train twice: Improving group robustness without training group information," in *International Conference on Machine Learning*. PMLR, 2021, pp. 6781–6792.
- [7] C. Fifty, E. Amid, Z. Zhao, T. Yu, R. Anil, and C. Finn, "Efficiently identifying task groupings for multi-task learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 27 503–27 516, 2021.
- [8] M. Geng, X. Xie, S. Liu, J. Yu, S. Hu, X. Liu, and H. Meng, "Investigation of data augmentation techniques for disordered speech recognition," *Proc. Interspeech*, 2022.
- [9] C. Bhat, A. Panda, and H. Strik, "Improved asr performance for dysarthric speech using two-stage dataaugmentation," *Proc. Interspeech*, pp. 46–50, 2022.
- [10] Z. Jin, M. Geng, X. Xie, J. Yu, S. Liu, X. Liu, and H. Meng, "Adversarial data augmentation for disordered speech recognition," *Proc. Interspeech*, 2021.
- [11] Z. Jin, X. Xie, M. Geng, T. Wang, S. Hu, J. Deng, G. Li, and X. Liu, "Adversarial data augmentation using vae-gan for disordered speech recognition," *arXiv preprint arXiv:2211.01646*, 2022.
- [12] B. Vachhani, C. Bhat, and S. K. Kopparapu, "Data augmentation using healthy speech for dysarthric speech recognition," *Proc. Interspeech*, 2018.
- [13] M. K. Baskar, T. Herzig, D. Nguyen, M. Diez, T. Polzehl, L. Burget, J. Černocký *et al.*, "Speaker adaptation for wav2vec2 based dysarthric asr," *Proc. Interspeech*, 2022.
- [14] S. Hu, X. Xie, Z. Jin, M. Geng, Y. Wang, M. Cui, J. Deng, X. Liu, and H. Meng, "Exploring self-supervised pre-trained asr models for dysarthric and elderly speech recognition," *arXiv preprint arXiv:2302.14564*, 2023.
- [15] L. P. Violeta, W.-C. Huang, and T. Toda, "Investigating self-supervised pretraining frameworks for pathological speech recognition," *Proc. Interspeech*, 2022.
- [16] A. Hernandez, P. A. Pérez-Toro, E. Nöth, J. R. Orozco-Arroyave, A. Maier, and S. H. Yang, "Cross-lingual self-supervised speech representations for improved dysarthric speech recognition," *Proc. Interspeech*, 2022.
- [17] M. Geng, X. Xie, R. Su, J. Yu, Z. Ye, X. Liu, and H. Meng, "On-the-fly feature based speaker adaptation for dysarthric and elderly speech recognition," *arXiv preprint arXiv:2203.14593*, 2022.
- [18] M. Geng, S. Liu, J. Yu, X. Xie, S. Hu, Z. Ye, Z. Jin, X. Liu, and H. Meng, "Spectro-temporal deep features for disordered speech assessment and recognition," *Proc. Interspeech*, 2021.
- [19] M. Geng, X. Xie, Z. Ye, T. Wang, G. Li, S. Hu, X. Liu, and H. Meng, "Speaker adaptation using spectro-temporal deep features for dysarthric and elderly speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 2597–2611, 2022.
- [20] 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.
- [21] T. Sainburg, M. Thielk, and T. Q. Gentner, "Finding, visualizing, and quantifying latent structure across diverse animal vocal repertoires," *PLoS computational biology*, vol. 16, no. 10, p. e1008228, 2020.
- [22] T. Sainburg, "timsainb/noisereducer: v1.0," Jun. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.3243139>
- [23] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *arXiv preprint arXiv:2212.04356*, 2022.
- [24] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Huggingface's transformers: State-of-the-art natural language processing," *arXiv preprint arXiv:1910.03771*, 2019.
- [25] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2019.