



Speaker adaptation for Wav2vec2 based dysarthric ASR

Murali Karthick Baskar^Φ, Tim Herzig[†], Diana Nguyen[†], Mireia Diez^Φ,
Tim Polzehl[†], Lukáš Burget^Φ and Jan “Honza” Černocký^Φ

^ΦBrno University of Technology, [†]Technische Universität Berlin

baskar@fit.vutbr.cz

Abstract

Dysarthric speech recognition has posed major challenges due to lack of training data and heavy mismatch in speaker characteristics. Recent ASR systems have benefited from readily available pretrained models such as wav2vec2 to improve the recognition performance. Speaker adaptation using fMLLR and xvectors have provided major gains for dysarthric speech with very little adaptation data. However, integration of wav2vec2 with fMLLR features or xvectors during wav2vec2 finetuning is yet to be explored. In this work, we propose a simple adaptation network for fine-tuning wav2vec2 using fMLLR features. The adaptation network is also flexible to handle other speaker adaptive features such as xvectors. Experimental analysis show steady improvements using our proposed approach across all impairment severity levels and attains 57.72% WER for high severity in UASpeech dataset. We also performed experiments on German dataset to substantiate the consistency of our proposed approach across diverse domains.

Index Terms: Dysarthria, self-supervision, ASR, wav2vec2, fMLLR, xvectors, speaker adaptation

1. Introduction

Automatic speech recognition (ASR) for dysarthric speech is hard due to two major constraints: 1) lack of sufficient data; and 2) unique speech characteristics of each speaker leading to domain mismatch. Data pooling [1] from multiple dysarthric speech datasets is one possible solution to handle the limited data issue. Data augmentation, GAN generated synthetic data [2, 3] and speech or volume perturbation [4] have also been used to address the lack of sufficient data to build dysarthric ASR systems. Recent works [5] addressed the mismatch issue [6] using domain adaptation and in [7], unsupervised domain adaptation from source to target domains is achieved using adversarial training.

Finetuning a pre-trained model using supervised dysarthric speech is a widely explored technique and has shown to tackle both issues simultaneously. In addition to naive finetuning, studies have shown better gains by using auxiliary speaker information. For instance, neural network based [8, 9] dysarthric ASR have shown that feature space maximum likelihood linear regression (fMLLR) is a relatively superior and simple speaker normalization approach. Authors in [10] have shown that xvectors act as complementary features to reduce speaker variability.

The advent of end-to-end (E2E) ASR systems such as Transformer based ASR outperformed the hybrid systems by a significant margin. E2E has been recently applied in [10] for dysarthric speech in combination with adaptation such as LHUC and xvectors. The authors in [11] use unsupervised bottleneck features in conjunction with fMLLR features for improving dysarthric speech recognition. Spectro-temporal deep features were used in [12] to remove the domain variance caused by articulatory

disfluencies and perform effective speaker adaptation using multiple dysarthric datasets. In [13], the authors used variational autoencoder to reduce the dysarthric speech variability as an enhancement to LHUC based model adaptation. A Bayesian gated DNN architecture [14] was found to perform better acoustic and pitch feature integration for improving robustness in dysarthric speech recognition. E2E based adaptation uses meta-learning approaches to adapt the parameters to unseen dysarthric speakers.

However, the recently proposed wav2vec2 [15] model has shown to improve E2E systems. In this work, we attempt to finetune a wav2vec2 model in conjunction with features such as fMLLR and xvectors within a wav2vec2 framework.

In this work, we propose a simple yet effective adaptation (finetuning) strategy named “wav2vec2 adapter” to incorporate wav2vec2 with speaker adaptation. Here, we attempt to finetune the wav2vec2 by feeding speaker information as auxiliary features during fine-tuning to efficiently finetune the wav2vec2 model parameters. An adapter network containing a bottleneck layer is instilled into the context encoder network of wav2vec2 model to integrate the auxiliary features and wav2vec2 outputs. The adapter module is based on the recent works [16, 17] on machine translation to fine-tune pre-trained transformer models. The primary contributions of this work are as follows:

- We hypothesize that our adapter approach is independent of the pretraining data domain used in wav2vec2 model and empirically verify it.
- Wav2vec2 adapter is shown to be flexible with auxiliary features such as fMLLR and xvectors.
- Effect of speaker dependent and independent finetuning on wav2vec2 models are shown by comparing it with hybrid systems.
- Crosslingual experiments with English (EN) and German (DE) are conducted using wav2vec2 adapter framework.

UASpeech dataset is used to conduct our experiments and results are calculated for different severity levels. Initially, we show the importance of wav2vec2 for E2E and its comparison with hybrid systems. Experimental analysis are further conducted to substantiate the impact of our proposed wav2vec2-adapter using fMLLR and xvectors separately and jointly.

2. Datasets used

UASpeech dataset is used in this work. It comprises 13 healthy control speakers and 15 dysarthric speakers. The recordings consists of isolated speech from each speaker uttering 765 words. The vocabulary includes 455 distinct words with 10 digits, 26 radio alphabets, 19 computer commands, 100 common words and 300 uncommon words. Speakers were divided into four different categories based on the severity of the condition, namely high (H), mid (M), low (L) and very low (VL).

German dataset is created from audio recordings of conversations between dysarthric patients and medical professionals. It contains a total of 21 dysarthric speakers, 6 females and 15 males. The corpus includes utterance level speech intelligibility ratings assessed by five naive listeners. The assessors assigned a subjective intelligibility assessment grade based on a five point scale, where "0" is assigned to clear speech and "4" being unintelligible. The intelligibility score is used to evaluate the severity of the impairment. Based on this, the speakers were divided into four different levels of severity (high: score 3 & 4, medium: score 2, low: score 1 and very low: score 0).

Table 1: Total amount of data in the German dataset for each dysarthric severity degree.

Severity	Duration (MM:SS)	Speakers
Very low	06:42	F03, M14, M17
Low	40:04	F02, F05, F06, M01, M09, M12, M20
Medium	17:19	M01, M03, M06, M10, M15
High	19:21	F01, F04, M02, M04, M08, M16

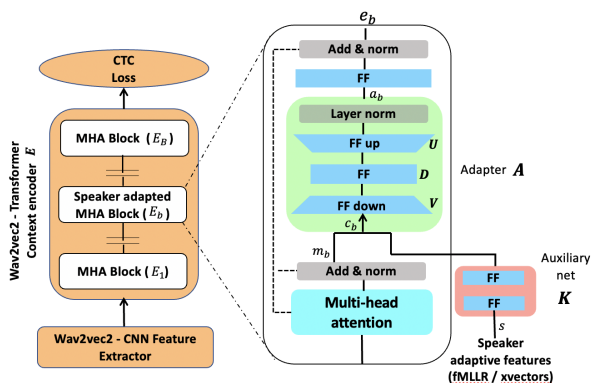


Figure 1: Adaptation framework inside a pre-trained wav2vec2 model using auxiliary speaker features. The green block denotes our proposed adapter module to inherit speaker information from the auxiliary net (pink block) into wav2vec2 for performing speaker adaptation.

3. Wav2vec2 Adapter

3.1. Preliminaries

The wav2vec2 are readily available pre-trained models trained with huge amounts of unsupervised data. Let \mathbf{X} be the raw input waveform. \mathbf{X} is sent to the feature extractor to obtain the encoded representations. The feature extractor contains convolutional layers and performs subsampling at a factor of 4 and directs its outputs to two parallel modules named masking component and quantizer. The masked inputs are passed into the context encoder network \mathbf{E} which inherits the transformer architecture. The encoder network \mathbf{E} contains B multi-head attention (MHA) blocks. Each MHA block E_b contains a MHA layer with layer norm followed by linear layer with layer norm and residual connections inbetween. The quantizer contains Gumbel-softmax activation which converts the unmasked inputs to quantized representations which act as labels for computation of contrastive loss training objective. In this work, we focus on adapting the parameters in the context encoder network \mathbf{E} by incorporating an additional network within a block E_b to normalize the dysarthric speaker characteristics.

3.2. Proposed approach

3.2.1. Motivation

Wav2vec2 pretrained models have shown drastic improvements in speech recognition performance. Our empirical study using UASpeech on finetuning a LF-MMI [18] and a wav2vec2 model to a particular speaker showed that LF-MMI improved in performance, while wav2vec2 did not converge. Besides, LF-MMI also showed further gains with fMLLR features, but speaker dependent finetuning is hard with wav2vec2. We hypothesize that the integration with fMLLR features can reap major benefits. We also hypothesize that incorporation of xvectors can act complementary to fMLLR features.

3.2.2. Model architecture

Our proposed wav2vec2 adapter framework addresses the limitations of fine-tuning wav2vec2 in a speaker dependent manner and can be used to effectively integrate auxiliary features. The framework comprises two components: 1) an adapter; and 2) an auxiliary net. The adapter is based on the recently proposed architecture [16] with layer normalized bottleneck layer and residual connection. Figure 1 shows the detailed block diagram view of our proposed framework. Here, we integrate the speaker adaptive features s through the auxiliary net (pink block) \mathbf{K} and integrate it with the adapter module \mathbf{A} present in the green block. The adapter \mathbf{A} is placed after the multi-head attention in a particular block E_b in the wav2vec2 model. It first projects down the hidden size h of the MHA block output, m_b to d dimensions using a feed-forward linear layer \mathbf{V} . This is followed by another d dimensional linear layer \mathbf{D} and projected back to h using another linear layer \mathbf{U} . $ReLU$ activation is present after every linear layer in the architecture. The adapter a_b output inside a particular block E_b , is defined as:

$$a_b = \mathbf{U}(\mathbf{D}(\mathbf{V}(c_b))) \quad (1)$$

$$\text{where, } c_b = \text{Append}((\mathbf{K}(s)), m_b) \quad (2)$$

Here,

$$s = \{f_1, f_2, \dots, f_T\} \in \mathbf{F}; \text{ for fMLLR features} \quad (3)$$

$$s = \{g_1, g_2, \dots, g_T\} \in \mathbf{G}; \text{ for xvectors} \quad (4)$$

f_t denotes a fMLLR feature at time t from a sequence \mathbf{F} and xvectors G is computed for each utterance u to get g_u and is repeated for all time steps T to get $g_{1:T} = g_u$. Unlike the adapter net in machine translation [16] which uses an adapter after every block, we only use a single adapter module after a particular block in wav2vec2. This setup helps to avoid drastic increase in parameters. We also use an auxiliary net and integrate it to the input of the adapter for instilling speaker adaptation during finetuning. Our proposed wav2vec2 adapter is flexible to integrate any type of auxiliary information to perform speaker normalization.

3.2.3. Auxiliary Features

In this work, we explore the wav2vec2 adapter module using supervised speaker adaptive features "fMLLR" and unsupervised speaker vectors "xvectors". These are described as follows:

fMLLR: The fMLLR features are obtained from the MLLR transforms trained with supervised adaptation data. Conventional fMLLR transformation is applied over MFCC or filterbank features. In our work, we extract input features for computing fMLLR transforms from the output (1024 dims) of the final

block E_B in figure 1 by forward propagating the raw waveforms as wav2vec2 input. We arrived to the usage of wav2vec2 based features for fMLLR computation and its integration at the B^{th} MHA block empirically. The collected 1024 dim. wav2vec2 features are transformed using LDA followed by GMM-HMM based training. Kaldi recipe¹ is used for LDA+MLLT+SAT training and speaker dependent fMLLR transforms computation to get 40 dimensional fMLLR features F .

Xvectors: Xvectors G are extracted from MFCC features using a separately trained xvector extractor. Since the xvectors are not directly extracted from wav2vec2 features, we integrate the xvectors with the output of the intermediate block i.e. E_b in figure 1, here $b = 2$ is decided empirically (see section 5.2). To substantiate our results, we follow a similar strategy to the one presented in [19], where the authors integrated ivectors for neural network adaptation in [19].

Finetuning with wav2vec2 adapter involves two stages: 1) the model parameters within the green and pink block are initially updated by fixing the rest of parameters. This first update helps the wav2vec2 encoder establish connections with the auxiliary information. 2) The whole model is then finetuned till convergence. Training configuration is explained in section 4.

4. Experimental setup

Architectures and Initialization: Kaldi [20] is used to train our baseline DNN-HMM, LF-MMI hybrid systems and the fMLLR feature extraction. Transformer and Wav2vec2 systems are built using ESPnet toolkit [21]. DNN-HMM contains 7 hidden layers each with 2048 dimensions. LF-MMI is built with 13 TDNN layers trained using sequence-level objective with data augmentations and ivectors. Our implementations will be available² with updates.

Transformer model contains 24 encoder blocks and 6 decoder layers each with 1024 dimensions. The transformer is initially pretrained with 960 hours of Librispeech data as we found random initialization did not lead to convergence with dysarthric data due to increased model complexity. Wav2vec2 follows similar architecture as the transformer and in this work we use two types of models: 1) wav2vec2 pretrained with 60k hours of Libri-light data and finetuned with 960 hours of Librispeech data³. Experiments on UASpeech are conducted using this wav2vec2 model. 2) XLSR wav2vec2 pretrained with 53 multilingual languages ($\approx 56k$ hours from the MLS, CommonVoice and Babel corpora) as defined in [22]. 2k hours of German data from CommonVoice is later used to finetune the XLSR⁴.

Training configuration: The wav2vec2 model is finetuned by optimizing using Adam with weight decay fix [23] with the parameters: $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-8$ and L_2 weight decay of 0. We set an initial learning rate of $1e-4$. The learning rate is warmed up over the first 500 steps then linearly decayed. We choose a total number of training epochs of 30. The global norm of the gradient is clipped at 1.0.

xvector extraction: 512 dimensional xvectors are extracted from a multilingual (34 languages) extractor trained with 1.65M recordings from 78k speakers using CNN-TDNN architecture. The detailed description of xvector extraction are in [24].

¹<https://github.com/kaldi-asr/kaldi/blob/master/egs/wsj/s5/run.sh#L211>

²<https://github.com/creatorscan/Dysarthric-ASR>

³<https://huggingface.co/facebook/wav2vec2-large-960h-lv60>

⁴<https://huggingface.co/facebook/wav2vec2-large-xlsr-53-german>

5. Results and discussion

We initially perform an experimental study on different model architectures such as DNN-HMM, LF-MMI and Transformer with fMLLR and model adaptation. The model adaptation (MA) is done by fine-tuning all the model parameters of DNN-HMM or LF-MMI in two stages: During stage 1, the model is trained with all supervised dysarthric data independent of speakers. The trained model is adapted to each speaker at stage 2. In case of Transformer MA, the model parameters are initialized with 960hrs of Librispeech, which is followed by speaker independent re-training with dysarthric data. We found randomly initialized Transformer did not result in convergence. Speaker dependent training also degraded the Transformer model performance which is in contrary to Hybrid models.

5.1. Effect of fMLLR features

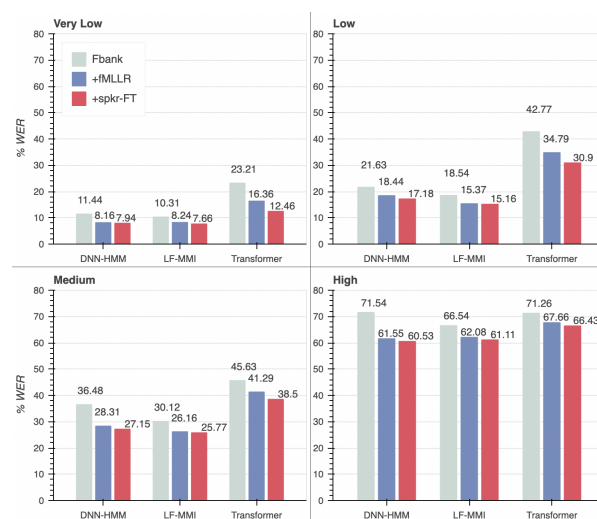


Figure 2: Performance comparison of fMLLR features with speaker level finetuning (spkr-FT) across different model architectures such as DNN-HMM, LF-MMI and Transformers

DNN-HMM, LF-MMI and Transformers are initially trained with fMLLR features of all speakers and later finetuned to each speaker. Figure 2 compares the models finetuned with and without fMLLR features. Among hybrid models, LF-MMI outperforms DNN-HMM in all cases. Transformer based ASR did not show any improvement when compared to either of the hybrid systems. Nevertheless, finetuning with fMLLR features has consistently outperformed the baseline filterbank based systems.

To improve the capability of Transformer, wav2vec2 model is used on behalf of transformer. Results are in table 2. Unlike the Transformer system, Wav2vec2 finetuned with dysarthric data shows comparable performance to LF-MMI. Finetuning wav2vec2 with our adapter module using fMLLR features significantly outperforms our best LF-MMI system.

5.2. Optimal MHA block for xvector Integration

Figure 3 shows the analysis to find the optimal block inside wav2vec2 encoder to integrate xvectors. Figure 3 shows a clear trend that feeding xvectors at higher level blocks resulted in performance degradation in terms of average %WER across all speakers. Therefore, based on empirical evidence, we select the second block E_2 to be used as an adapter module to integrate

Table 2: WER% results for our wav2vec2 adapter module, finetuned (FT) with fMLLR features, with xvectors and with their combination

Model	Adaptation	VL	L	M	H
LF-MMI	fMLLR+Spkr-FT	7.66	15.16	25.77	61.11
Wav2vec2	FT	8.54	16.67	33.70	64.78
	fMLLR+FT	7.42	12.07	23.11	58.48
	xvectors+FT	7.11	13.42	23.52	59.52
	xvectors+fMLLR+FT	6.83	12.82	22.46	57.72

xvectors (unlike fMLLR features which are fed at the last block).




Figure 3: Performance comparison of integrating xvectors at different blocks inside wav2vec2

Table 2 shows the importance of adapting the wav2vec2 using fMLLR and xvectors. fMLLR features shows clear wins compared to xvectors for all severities except the VL severity. Adapting the wav2vec2 with two adapters, one placed at first block fed with xvectors and the second adapter in the last block fed with fMLLR features show complementarity with each other by attaining consistent gains.

Table 3: Crosslingual experiments on English (EN) and German (DE) with model adaptation

Type	Training	Finetuning	Evaluation	Avg. %WER
Baseline	EN	-	EN	24.21
	DE	-	DE	28.43
Crosslingual	EN+DE	-	EN	24.25
		-	DE	27.32
	EN	EN	EN	23.76
		DE	DE	25.93

5.3. Cross-lingual experiments on English and German

A crosslingual (XLSR) model is initially built by combining EN from UASpeech and DE from our German corpus. The output layer contains 75 dimensions by combining 40 phonemes from English and 35 phonemes from German along with their corresponding language tag. The output layer is connected to the pre-trained XLSR model followed by fine-tuning with both EN and DE together and the resulting model is evaluated for EN and DE. The results in table 3 show a 4.2% relative improvement for DE over the baseline system but no improvement for EN. This is because DE has very little training examples compared to EN, which makes the crosslingual training beneficial for DE. Further gains are obtained by finetuning to each language independently and the results are in table 3.

5.4. Comparison across Literature

Among several works [26, 10, 25, 4] on UASpeech, authors in CUHK-2019 [25], showed that lip features from video helps

Table 4: Performance of wav2vec2 adapter, CUHK system on 3 impaired speakers in UASpeech

ID	F04 (Low)	M12 (High)	M14 (VeryLow)
Dysarthria	Athetoid CP	Mixed	Spastic CP
CUHK [25]	24.0	54.0	15.6
Ours	22.5	49.9	12.9

to improve the hybrid DNN system. In CUHK-2021, [4] performed neural architecture search (NAS) and combined it with data augmentation (DA), video embedding (AV) and LHUC based adaptation. In 2022, [12] showed that hybrid DNN trained with spectro-temporal features and xvectors outperformed their Conformer system. Spectro-temporal features capture longer context for better normalization of speaker variability and hence improves over CUHK system. In table 4 we show how our proposed system behaves and improves over CUHK system [25, 10] across different types of dysarthric speakers. Table 5 shows the average %WER across all severity levels between existing best systems. Our proposed wav2vec2 adapter and its multilingual extension attains 3.35% and 5.15% relative improvement over SBE [12] respectively.

Table 5: A comparison between published works on UASpeech and our proposed work

Systems	Avg. %WER
CUHK-2021 NAS DNN+ DA + LHUC + AV [10]	25.21
DA + SBE + LHUC [12]	25.05
Wav2vec2 +fMLLR+xvectors (Ours)	22.83
Wav2vec2 (XLSR) +fMLLR+xvectors+Mult. (Ours)	22.48

6. Conclusion

In this work, we investigate the importance of speaker dependent auxiliary features such as fMLLR and xvectors for adapting wav2vec2 models for improving dysarthric speech recognition. We show that in contrast to hybrid systems, wav2vec2 did not improve by adapting model parameters based on each speaker. We proposed a wav2vec2 adapter module that inherits speaker features as auxiliary information to perform effective speaker normalization during finetuning. We showed that, using the adapter module, fMLLR and xvectors are complementary to each other, and proved the effectiveness of the approach outperforming existing SoTA on UASpeech dysarthric speech ASR. In our cross-lingual experiments, we also showed that combining English and German data for training, can further improve performance of our systems, proving useful in scenarios where little training examples exist for a particular language.

7. Acknowledgements

The work was supported by European Union’s Horizon 2020 project HumaneAI-Net, No. 761758, Czech National Science Foundation (GACR) project NEUREM3 No. 19-26934X, and Czech Ministry of Education, Youth and Sports project no. LTAIN19087 ”Multi-linguality in speech technologies”. Computing on IT4I supercomputer was supported by the Czech Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations project ”e-Infrastructure CZ-LM2018140”.

8. References

- [1] S. R., M. Ramasubba Reddy, and S. Umesh, "Improved acoustic modeling for automatic dysarthric speech recognition," in *2015 Twenty First National Conference on Communications (NCC)*, pp. 1–6, 2015.
- [2] Z. Jin, M. Geng, X. Xie, J. Yu, S. Liu, X. Liu, and H. Meng, "Adversarial data augmentation for disordered speech recognition," *arXiv preprint arXiv:2108.00899*, 2021.
- [3] B. Vachhani, C. Bhat, and S. K. Kopparapu, "Data augmentation using healthy speech for dysarthric speech recognition.," in *Inter-speech*, pp. 471–475, 2018.
- [4] S. Liu, S. Hu, X. Xie, M. Geng, M. Cui, J. Yu, X. Liu, and H. M. Meng, "Recent progress in the cuhk dysarthric speech recognition system," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2021.
- [5] F. Xiong, J. Barker, Z. Yue, and H. Christensen, "Source domain data selection for improved transfer learning targeting dysarthric speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7424–7428, IEEE, 2020.
- [6] I. Divyasree and K. Selvamani, "Dad: Domain adversarial defense system against ddos attacks in cloud," *IEEE Transactions on Network and Service Management*, 2021.
- [7] D. Wang, L. Deng, Y. T. Yeung, X. Chen, X. Liu, and H. Meng, "Unsupervised domain adaptation for dysarthric speech detection via domain adversarial training and mutual information minimization," *arXiv preprint arXiv:2106.10127*, 2021.
- [8] N. M. Joy and S. Umesh, "Improving acoustic models in torgo dysarthric speech database," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 26, no. 3, pp. 637–645, 2018.
- [9] F. Xiong, J. Barker, and H. Christensen, "Deep learning of articulatory-based representations and applications for improving dysarthric speech recognition," in *Speech Communication; 13th ITG-Symposium*, pp. 1–5, VDE, 2018.
- [10] 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)*, pp. 1–5, IEEE, 2021.
- [11] Z. Yue, H. Christensen, and J. Barker, "Autoencoder bottleneck features with multi-task optimisation for improved continuous dysarthric speech recognition," in *Proceedings of Interspeech 2020*, pp. 4581–4585, International Speech Communication Association (ISCA), 2020.
- [12] 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," *arXiv preprint arXiv:2202.10290*, 2022.
- [13] X. Xie, R. Ruzi, X. Liu, and L. Wang, "Variational auto-encoder based variability encoding for dysarthric speech recognition," *arXiv preprint arXiv:2201.09422*, 2022.
- [14] S. Liu, S. Hu, X. Liu, and H. Meng, "On the use of pitch features for disordered speech recognition.," in *INTERSPEECH*, pp. 4130–4134, 2019.
- [15] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *arXiv preprint arXiv:2006.11477*, 2020.
- [16] J. Pfeiffer, I. Vulić, I. Gurevych, and S. Ruder, "Mad-x: An adapter-based framework for multi-task cross-lingual transfer," *arXiv preprint arXiv:2005.00052*, 2020.
- [17] J. Pfeiffer, A. Rücklé, C. Poth, A. Kamath, I. Vulić, S. Ruder, K. Cho, and I. Gurevych, "Adapterhub: A framework for adapting transformers," *arXiv preprint arXiv:2007.07779*, 2020.
- [18] D. Povey, V. Peddinti, D. Galvez, P. Ghahramani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi," *INTERSPEECH*, 2016.
- [19] M. Karafiát, M. K. Baskar, P. Matějka, K. Veselý, F. Grézl, and J. Černocký, "Multilingual BLSTM and speaker-specific vector adaptation in 2016 BUT Babel system," in *Spoken Language Technology Workshop (SLT), 2016 IEEE*, pp. 637–643, IEEE, 2016.
- [20] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, *et al.*, "The kaldi speech recognition toolkit," in *Automatic Speech Recognition and Understanding, 2011 IEEE Workshop on*, pp. 1–4, IEEE, 2011.
- [21] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, *et al.*, "Espnet: End-to-end speech processing toolkit," *arXiv preprint arXiv:1804.00015*, 2018.
- [22] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, "Unsupervised cross-lingual representation learning for speech recognition," *arXiv preprint arXiv:2006.13979*, 2020.
- [23] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [24] M. Karafiát, K. Veselý, J. Profant, J. Nytra, M. Hlaváček, T. Pavlíček, *et al.*, "Analysis of x-vectors for low-resource speech recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6998–7002, IEEE, 2021.
- [25] S. Hu, S. Liu, H. F. Chang, M. Geng, J. Chen, L. W. Chung, T. K. Hei, J. Yu, K. H. Wong, X. Liu, *et al.*, "The cuhk dysarthric speech recognition systems for english and cantonese.," in *INTER-SPEECH*, pp. 3669–3670, 2019.
- [26] G. Wang, A. Rosenberg, Z. Chen, Y. Zhang, B. Ramabhadran, Y. Wu, and P. Moreno, "Improving speech recognition using consistent predictions on synthesized speech," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7029–7033, IEEE, 2020.