

Investigation of Data Augmentation Techniques for Disordered Speech Recognition

Mengzhe Geng^{1*}, Xurong Xie^{1,2*}, Shansong Liu¹, Jianwei Yu¹, Shoukang Hu¹,
Xunying Liu¹, Helen Meng¹

¹The Chinese University of Hong Kong

²Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences

{mzgeng, sslui, jwyu, skhu, xyliu, hmmeng}@se.cuhk.edu.hk, xrxie@ee.cuhk.edu.hk

Abstract

Disordered speech recognition is a highly challenging task. The underlying neuro-motor conditions of people with speech disorders, often compounded with co-occurring physical disabilities, lead to the difficulty in collecting large quantities of speech required for system development. This paper investigates a set of data augmentation techniques for disordered speech recognition, including vocal tract length perturbation (VTLP), tempo perturbation and speed perturbation. Both normal and disordered speech were exploited in the augmentation process. Variability among impaired speakers in both the original and augmented data was modeled using learning hidden unit contributions (LHUC) based speaker adaptive training. The final speaker adapted system constructed using the UASpeech corpus and the best augmentation approach based on speed perturbation produced up to 2.92% absolute (9.3% relative) word error rate (WER) reduction over the baseline system without data augmentation, and gave an overall WER of 26.37% on the test set containing 16 dysarthric speakers.

Index Terms: Speech Disorders, Speech Recognition, Data Augmentation, Speaker Adaptive Training

1. Introduction

Speech disorders affect millions of people all over the world, severely degrading their quality of life. They may be caused by a range of medical conditions, such as cerebral palsy [1], Parkinson disease [2] and stroke [3]. Speech disorders are always associated with difficulties in controlling articulators responsible for speech, which creates a large mismatch against normal speech [4]. This makes the recognition of disordered speech a challenging task [5, 6]. Meanwhile, as speech disorders always come together with loss of abilities in physical motion, there is a popular demand of hands-free and speech-enabled assistive technologies to help such people [7, 8].

Resurgence of deep learning technologies in the past decade improved the performance of state-of-the-art automatic speech recognition systems [9, 10, 11, 12, 13, 14]. However, these systems are not directly usable by people with speech disorders due to the large mismatch against disordered speech. Hence, there has been growing research interest in developing deep neural network (DNN) based modeling techniques suitable for disordered speech [15, 16, 17, 18, 19, 20, 21]. The underlying neuro-motor conditions of people suffering from speech impairment, often compounded with co-occurring physical disabilities, lead to the difficulty in collecting large quantities of disordered speech required for automatic speech recognition (ASR)

system development [22, 23, 24]. In addition, varying speech disorder characteristics and levels of severity create a large variation among different dysarthric speakers. These factors combined impose limitations on the performance of current disordered speech recognition systems [17, 18, 19, 20, 21].

Data augmentation proves to be an effective method for dealing with data sparsity and improving the performance of DNN acoustic models for normal speech recognition [14, 25, 26, 27, 28, 29, 30], as it enables the neural networks to improve coverage and generalization. A wide range of augmentation techniques have been investigated, including cross domain adaptation [25], vocal tract length perturbation (VTLP) [26, 28, 29], spectral distortion [27], tempo perturbation [27, 29], stochastic feature mapping [28], speed perturbation [29] and end-to-end back translation [30]. Various spectrogram deformations were applied in [14].

In contrast, so far there is only limited research on data augmentation targeting disordered speech recognition. In [16], normal speech data was applied in the bottleneck feature extraction stage. In [31], data augmentation based on time- and tempo-stretching of normal speech was investigated. In [32], a voice conversion framework was proposed to convert normal speech to disordered speech. In [33], phonetic analysis was conducted to compute the speaker-dependent scaling factors. These were then used in tempo adjustment to augment the limited disordered speech data. There are two issues associated with previous research. First, previous work in this direction provided piece-wise solutions, lacking a systematic comparison between augmentation approaches for disordered speech recognition. Second, the main focus of previous research has been on adjusting normal speech towards disordered speech, while investigation on augmentation using existing disordered speech data has been very limited.

This paper presents a systematic investigation of different data augmentation techniques for disordered speech recognition. Augmented data was generated from two different sources: i) modification of healthy speech to disordered speech, and ii) perturbation of existing disordered speech. Three data augmentation techniques were used, including vocal tract length perturbation (VTLP), tempo perturbation and speed perturbation. Tempo perturbation stretches the length of the utterance but keeps the shape of spectral envelope unchanged [34], while speed perturbation resamples signals in time domain [29]. Speaker-level perturbation factors were obtained based on phonetic analysis described in [33] and then used on normal speech. Global perturbation factors were used on disordered speech. Performance evaluation was conducted on a speaker independent DNN system implemented using an extended version of the Kaldi toolkit [35] for the Universal Access Speech (UASpeech)

*Equal contribution was made between the first two authors.

database. Learning hidden unit contributions (LHUC) [36] based speaker adaptive training (SAT) was further applied to model the large variability among disordered speakers in both the original and augmented data.

The main contributions of this paper are summarized below. To the best of our knowledge, this is the first work to systematically investigate different data augmentation techniques for disordered speech recognition. Both normal and disordered speech were exploited in the augmentation process. Out of the three augmentation techniques, applying speed perturbation to both normal and disordered speech was found to give the best performance. The system using the best augmentation approach based on speed perturbation produced 2.92% absolute (9.3% relative) word error rate (WER) reduction over the baseline system without data augmentation. The final speaker adaptive system using the augmented data gave an overall WER of 26.37% on the test set containing 16 UASpeech dysarthric speakers. As far as we know, this is the best performance reported so far on UASpeech.

The rest of this paper is organized as follows. A range of data augmentation techniques are presented in section 2. Section 3 describes the baseline multi-speaker adaptively trained DNN system architecture. Section 4 presents experiments and results on the UASpeech database. The last section concludes and discusses possible future works.

2. Data Augmentation

We investigate three data augmentation techniques for disordered speech recognition. Both disordered speech and normal speech were used in the augmentation process. The former applied global perturbation factors while the latter applied speaker dependent perturbation factors based on phonetic analysis.

2.1. VTLP Based Data Augmentation

The shape of vocal tract varies from speaker to speaker [37]. Vocal tract length normalization (VTLN) was proposed to counteract such variations [38, 39]. By applying VTLN in the reverse manner, vocal tract length perturbation (VTLP) adds variability to the speech data by simulating different vocal tract lengths [26]. Given a time-domain audio segment $x(t)$, we denote the corresponding frequency-domain representation as $X(f)$, which is the Fourier transform of $x(t)$. VTLP applies a perturbation factor α taken from a discrete set (e.g. $\{0.9, 1.1\}$) along the frequency axis of $X(f)$. The output $Y(f)$ is given as:

$$Y(f) = X(\alpha f) \quad (1)$$

In this way, VTLP perturbs the spectral envelope of the audio segment while keeping the audio duration unchanged.

2.2. Tempo Perturbation Based Data Augmentation

Tempo perturbation stretches the duration of the audio signal $x(t)$ while leaving the shape of its spectral envelope untouched [27, 29]. This is achieved by first decomposing the time-domain audio segment $x(t)$ into short analysis blocks and then relocating these blocks along the time axis to construct the perturbed output $y(t)$ [40]. A well known algorithm in this area is waveform similarity overlap-add (WSOLA), which makes $y(t)$ share the maximal similarity with $x(t)$ by finding the optimal position of each analysis block iteratively [34, 40].

Suppose the time-domain audio signal $x(t)$ is decomposed into short analysis blocks $\tilde{x}_m(r)$. These blocks are equally spaced along the time axis by H_a (the analysis hopsize). Given a perturbation factor α , the synthesis blocks $\tilde{y}_m(r)$ are relocated

along the time axis by H_s (the synthesis hopsize) given as:

$$H_s = \alpha \cdot H_a \quad (2)$$

WSOLA takes an iterative approach to update the positions of analysis blocks. For example, the center of $\tilde{x}_m(r)$ is shifted by $\Delta_m \in [-\Delta_{max}, \Delta_{max}]$ along the time axis, where the optimal value of Δ_m is obtained by maximizing the cross-correlation between $\tilde{x}_m(r)$ and $\tilde{x}_{m-1}(r)$. This ensures that the periodic structures of the adjusted analysis frame are optimally aligned with structures of the previously copied synthesis frame in the overlapping region when both frames use the synthesis hopsize H_s . The hann window function $w(r)$ is then applied on the adjusted analysis block to compute the synthesis block $\tilde{y}_m(r)$ [34, 40]. After finishing all iterations, the synthesis frames are processed in order to reconstruct the actual time-scale modified output signal $y(t)$ in a similar manner as conventional OLA [41]. In this way, $y(t)$ keeps the same spectral envelope shape as $x(t)$, but has a different length.

2.3. Speed Perturbation Based Data Augmentation

Speed perturbation resamples the audio signal in time domain [29]. Given an audio segment $x(t)$, a perturbation factor α is applied along the time axis and gives the output $y(t)$ as:

$$y(t) = x(\alpha t) \quad (3)$$

In frequency domain, this is equivalent to the following change:

$$X(f) \rightarrow \frac{1}{\alpha} X\left(\frac{1}{\alpha} f\right) \quad (4)$$

where $X(f)$ and $\frac{1}{\alpha} X\left(\frac{1}{\alpha} f\right)$ represent the Fourier transform of $x(t)$ and $y(t)$ respectively. In this way, speed perturbation leads to both change in the audio duration and perturbation in the spectral envelope [29, 42]. Table 1 summarizes the implement domain and effects of VTLP, tempo perturbation (denoted as ‘‘tempo’’) and speed perturbation (denoted as ‘‘speed’’).

Table 1: Comparison of implement domain and effects of VTLP, tempo perturbation and speed perturbation. ‘‘✓’’ indicates that change occurs after perturbation.

	VTLP	tempo	speed
Implement Domain	$X(f)$	$x(t)$	$x(t)$
Signal Duration	no change	✓	✓
Spectral Envelope	✓	no change	✓

2.4. Speaker Dependent Perturbation Factor Estimation

Given the large mismatch in speaking rate between normal and disordered speakers and high variability among disordered speakers, speaker dependent perturbation factors were applied when modifying normal speech to disordered speech. The factors were obtained based on phonetic analysis described in [33]. We performed force alignment using a GMM-HMM system based on the HTK toolkit [43] to get frame-level phoneme alignments, and then calculated the average phoneme duration of each speaker (denoted as l_{C_i} for control speaker C_i and l_{D_j} for dysarthric speaker D_j). For dysarthric speaker D_j , we took the average of l_{C_i} (denoted as $l_{\bar{C}}$) as the reference to compute the speaker dependent perturbation factor F_{D_j} , given as:

$$F_{D_j} = \frac{l_{\bar{C}}}{l_{D_j}} \quad (5)$$

F_{D_j} was then used as the perturbation factor when modifying normal speech to the speech of dysarthric speaker D_j .




Figure 1: Architecture of the hybrid DNN system. The dashed line indicates the connecting component is an optional block. Retaining connection (a) leads to the speaker-independent baseline system, while retaining connection (b) leads to the system with LHUC SAT.

3. Baseline ASR System Description

This section describes the baseline ASR system in terms of the hybrid DNN architecture and learning hidden unit contributions (LHUC) [36] based speaker adaptive training.

3.1. Hybrid Deep Neural Network Architecture

The structure of the hybrid deep neural network (DNN) acoustic model contains seven hidden layers. Each hidden layer contains a basic set of neural operations performed in sequence, i.e. affine transformation (in green), rectified linear unit (ReLU) activation (in yellow) and batch normalization (in orange). Apart from this, linear bottleneck projections (in light green) are applied to the inputs of the intermediate five hidden layers to reduce the number of network parameters. Dropout operations (in grey) are applied to the outputs of the first six hidden layers to prevent over-fitting. Softmax activation (in dark green) is used in the output layer. To speed up the training process and circumvent the vanishing gradient problem, two skip connections are used to connect the output of the first layer to the third layer and the output of the fourth layer to the sixth layer respectively. Figure 1 illustrates the architecture of the hybrid DNN system.

Multi-task learning (MTL) [44] was used to train the hybrid DNN system shown in Figure 1. The labels for the two tasks are based on frame-level tied triphone states and monophone alignments respectively. The alignments were obtained from a GMM-HMM system implemented using the HTK toolkit [43]. Incorporating frame-level monophone alignments in the labels reduces the risk of over-fitting to unreliable frame-level triphone states computed from disordered speech. The loss function of the multi-task learning is as follows:

$$\mathcal{L}_{MTL} = \lambda \cdot \mathcal{L}_{tristate} + (1 - \lambda) \cdot \mathcal{L}_{mono} \quad (6)$$

where $\mathcal{L}_{tristate}$ is the cross-entropy loss for the tied triphone state task, \mathcal{L}_{mono} is the cross-entropy loss for the monophone task, and $0 \leq \lambda \leq 1$ is a tunable task weight parameter.

3.2. LHUC Based Speaker Adaptive Training

To model the large variability among disordered speakers, learning hidden unit contributions (LHUC) [36] based speaker adaptive training (SAT) was used. Speaker-level LHUC scaling factors were applied to the ReLU activation output in the first layer, as shown by connection (b) in Figure 1. Supervised estimation of the LHUC factors was performed during the training stage, where the LHUC factors were updated once per mini-batch together with the network parameters. Unsupervised LHUC adaptation was performed during the test stage, where the LHUC factors were updated once per utterance.

4. Experiments and Results

4.1. Task Description

The UASpeech [22] is an isolated word recognition task consisting of 16 dysarthric and 13 control speakers. The speech materials contain 155 common words and 300 uncommon words. There are 3 blocks per speaker, each containing all 155 common words and one third of the uncommon words. We treated block 1 and block 3 of all 29 speakers as training set, and block 2 of the 16 dysarthric speakers as test set. Silence stripping was performed using a GMM-HMM system to remove redundant silence in the recordings, as described in our previous work [18]. For the baseline system without data augmentation, the training set contained 99195 utterances (~ 30.6 hours) and the test set contained 26520 utterances (~ 9 hours).

4.2. Experiment Setup

We implemented the techniques discussed in section 2 to perform data augmentation on the training set and left the test set untouched. The HTK toolkit [43] was used for VTLP. The tempo command of the Sox toolkit [45] was used for tempo perturbation, which is based on WSOLA. The speed command of Sox was used for speed perturbation. Following [29], three sets of global perturbation factors were applied on the disordered speech: $\{0.9, 1.1\}$, $\{0.9, 0.95, 1.05, 1.1\}$ and $\{0.85, 0.9, 0.95, 1.05, 1.1, 1.15\}$. Speaker dependent perturbation factors were applied when modifying normal speech to target disorder speech, as discussed in section 2.4. The speaker-level scaling values are shown in Figure 2.




Figure 2: Speaker-level scaling factors for modifying normal speech to target disordered speech in the order of intelligibility [15, 22]. A factor $F < 1$ slows down the normal speech while $F > 1$ speeds up the normal speech.

In our experiments, the hybrid DNN acoustic model was implemented using the Kaldi toolkit [35]. A 9-frame context window was used in the system. The inputs were 80-dimension filter bank (FBK) + Δ features. The first six hidden layers contain 2000 neurons each, while the dimension of the linear bottleneck projections is 200 and dropout rate is 20%. The seventh hidden layer contains 100 neurons. The system was trained by

Table 2: Performance on the 16 UASpeech dysarthric speakers of different data augmentation approaches applied on the training set. Here “CTL” / “DYS” stands for normal / disordered speech. “tempo” / “speed” stands for tempo / speed perturbation. “2x”, “4x” and “6x” refer to the amount of augmented data. “Very low”, “Low”, “Mid” and “High” refer to intelligibility of the DYS groups.

Sys. ID	Data Augmentation			# Hours	LHUC SAT	WER%				
	Method	CTL	DYS			Very low	Low	Mid	High	Overall
1	NA			30.6	✗	69.82	32.61	24.53	10.40	31.45
2	NA			30.6	✓	64.39	29.88	20.27	8.95	28.29
3	VTLP	1x	-	48.0	✗	68.68	31.84	22.71	9.48	30.35
4	Tempo	1x	-	52.2	✗	70.71	32.78	25.12	10.32	31.77
5	Speed	1x	-	52.2	✗	67.52	31.55	21.96	9.57	29.92
6	VTLP	-	2x	65.5	✗	69.98	30.08	21.39	9.65	29.97
7	Tempo	-	2x	65.9	✗	69.32	31.75	23.94	10.07	30.90
8	Speed	-	2x	65.9	✗	68.43	29.60	21.37	10.44	29.79
9	Speed	-	4x	100.9	✗	67.20	29.86	21.45	10.04	29.47
10	Speed	-	6x	136.7	✗	67.15	30.07	21.25	10.17	29.52
11	Speed	2x	2x	130.1	✗	66.45	28.95	20.37	9.62	28.73
12	Speed	4x	4x	207.5	✗	66.26	28.60	19.90	9.68	28.53
13	Speed	2x	2x	130.1	✓	62.50	27.26	18.41	8.04	26.55
14	Speed	4x	4x	207.5	✓	62.44	27.55	17.35	7.93	26.37

back-propagation based on RMSProp optimizer. For multi-task learning, the task weight parameter λ was set as 0.5, giving the same weight to both tied triphone state and monphone task. Following [15], a uniform language model was used in decoding.

4.3. Performance of Data Augmentation

Table 2 shows the performance of different data augmentation approaches applied on the training set. “2x”, “4x” and “6x” refer to the amount of augmented data. Several trends can be observed. **1)** Sys.3-5 show that for augmentation using normal speech, speed perturbation gives the best overall performance with an absolute WER reduction of 1.53% over the baseline (Sys.1). **2)** Sys.6-8 demonstrate that for augmentation using disordered speech, speed perturbation also gives the best overall performance. Therefore, we proceed with speed perturbation to generate more augmented data in the remaining experiments. **3)** Sys.8-10 indicate that perturbing the disordered speech by {0.9, 0.95, 1.05, 1.1}(4x) works better than {0.9, 1.1}(2x) or {0.85, 0.9, 0.95, 1.05, 1.1, 1.15}(6x). **4)** Sys.8-9,11-12 show that augmentation using both normal speech (with speaker dependent factors) and disordered speech (with global factors) can further improve the performance by an absolute WER reduction of 0.94% (Sys.12 over Sys.9), while increasing the amount of augmented data produces only marginal improvement (Sys.12 over Sys.11). The best speaker-independent system with data augmentation (Sys.12) produces **2.92%** absolute (9.3% relative) WER reduction over the speaker-independent baseline system without data augmentation (Sys.1).

4.4. Performance of LHUC SAT

LHUC based speaker adaptive training was further applied to model the variability among disordered speakers in both the original and augmented data (see Table 2, Sys.2,13-14). The augmented data was assigned to individual target dysarthric speakers for the speaker-level LHUC adaptation. This leads to a further 2.16% absolute WER reduction (Sys.14 over Sys.12), indicating that LHUC SAT contributes to modelling the large variability among disordered speakers. The final speaker adaptive system using the best augmentation approach gives an overall WER of 26.37% on the test set containing 16 UASpeech

dysarthric speakers (see Table 2, Sys.14). To the best of our knowledge, this corresponds to an absolute WER reduction of 1.51% over the best published system on UASpeech, as shown in Table 3.

Table 3: A comparison between published systems on UASpeech and our system. Here “DA” refers to data augmentation.

Systems	WER%
Sheffield-2013 Cross domain augmentation [16]	37.50
Sheffield-2015 Speaker adaptive training [17]	34.80
CUHK-2018 DNN System Combination [18]	30.60
Sheffield-2019 Kaldi TDNN + DA [33]	27.88
Speed perturb + LHUC SAT (Table 2, Sys.14)	26.37

5. Conclusions

This paper presents a systematic investigation of different data augmentation techniques for disordered speech recognition. It suggests that speed-perturbation based augmentation produces the largest improvement in system performance despite the huge mismatch between normal and disordered speech. Future research will focus on more powerful data augmentation techniques to cover other features of disordered speech, such as articulation imprecision, reduced intensity and disfluency.

6. Acknowledgement

This research is supported by Hong Kong Research Grants Council General Research Fund No. 14200218, Theme Based Research Scheme T45-407/19N and Shun Hing Institute of Advanced Engineering Project No. MMT-p1-19. The authors would like to thank Dr. Feifei Xiong for fruitful discussions.

7. References

- [1] 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.
- [2] S. Scott and F. Caird, “Speech therapy for Parkinson’s disease.”

- Journal of Neurology, Neurosurgery & Psychiatry*, vol. 46, no. 2, pp. 140–144, 1983.
- [3] P. Jerntorp and G. Berglund, “Stroke registry in Malmö, Sweden.” *Stroke*, vol. 23, no. 3, pp. 357–361, 1992.
 - [4] K. M. Rosen, R. D. Kent, A. L. Delaney *et al.*, “Parametric quantitative acoustic analysis of conversation produced by speakers with dysarthria and healthy speakers,” *Journal of Speech, Language, and Hearing Research*, 2006.
 - [5] R.-E. Hux, Karen, Joan, N. Manasse *et al.*, “Accuracy of three speech recognition systems: Case study of dysarthric speech,” *Augmentative and Alternative Communication*, vol. 16, no. 3, pp. 186–196, 2000.
 - [6] V. Young and A. Mihailidis, “Difficulties in automatic speech recognition of dysarthric speakers and implications for speech-based applications used by the elderly: A literature review,” *Assistive Technology*, vol. 22, no. 2, pp. 99–112, 2010.
 - [7] M. Fried-Oken, “Voice recognition device as a computer interface for motor and speech impaired people.” *Archives of physical medicine and rehabilitation*, vol. 66, no. 10, pp. 678–681, 1985.
 - [8] M. S. Hawley, S. P. Cunningham, P. D. Green *et al.*, “A voice-input voice-output communication aid for people with severe speech impairment,” *IEEE Transactions on neural systems and rehabilitation engineering*, vol. 21, no. 1, pp. 23–31, 2012.
 - [9] W. Xiong, L. Wu, F. Alleva *et al.*, “The Microsoft 2017 conversational speech recognition system,” in *ICASSP*. IEEE, 2018, pp. 5934–5938.
 - [10] G. Saon, G. Kurata, T. Sercu *et al.*, “English Conversational Telephone Speech Recognition by Humans and Machines,” in *INTERSPEECH*, 2017, pp. 132–136.
 - [11] S. Thomas, M. Suzuki, Y. Huang *et al.*, “English broadcast news speech recognition by humans and machines,” in *ICASSP*. IEEE, 2019, pp. 6455–6459.
 - [12] V. Manohar, D. Povey, and S. Khudanpur, “JHU Kaldi system for Arabic MGB-3 ASR challenge using diarization, audio-transcript alignment and transfer learning,” in *2017 IEEE ASRU Workshop*. IEEE, 2017, pp. 346–352.
 - [13] C. Lüscher, E. Beck, K. Irie *et al.*, “RWTH ASR Systems for LibriSpeech: Hybrid vs Attention,” in *INTERSPEECH*, 2019, pp. 231–235.
 - [14] D. S. Park, W. Chan, Y. Zhang *et al.*, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *INTERSPEECH*, 2019, pp. 2613–2617.
 - [15] H. Christensen, S. Cunningham, C. Fox *et al.*, “A comparative study of adaptive, automatic recognition of disordered speech,” in *INTERSPEECH*, 2012, pp. 1776–1779.
 - [16] H. Christensen, M. Aniol, P. Bell *et al.*, “Combining in-domain and out-of-domain speech data for automatic recognition of disordered speech,” in *INTERSPEECH*, 2013, pp. 3642–3645.
 - [17] S. Sehgal and S. Cunningham, “Model adaptation and adaptive training for the recognition of dysarthric speech,” in *Proceedings of SLPAT 2015*, 2015, pp. 65–71.
 - [18] J. Yu, X. Xie, S. Liu *et al.*, “Development of the CUHK Dysarthric Speech Recognition System for the UASpeech Corpus,” in *INTERSPEECH*, 2018, pp. 2938–2942.
 - [19] S. Liu, S. Hu, Y. Wang *et al.*, “Exploiting Visual Features using Bayesian Gated Neural Networks for Disordered Speech Recognition,” in *INTERSPEECH*, 2019, pp. 4120–4124.
 - [20] S. Liu, S. Hu, X. Liu *et al.*, “On the Use of Pitch Features for Disordered Speech Recognition,” in *INTERSPEECH*, 2019, pp. 4130–4134.
 - [21] S. Hu, S. Liu, H. F. Chang *et al.*, “The CUHK Dysarthric Speech Recognition Systems for English and Cantonese,” in *INTERSPEECH*, 2019, pp. 3669–3670.
 - [22] H. Kim, M. Hasegawa-Johnson, A. Perlman *et al.*, “Dysarthric speech database for universal access research,” in *INTERSPEECH*, 2008.
 - [23] D.-L. Choi, B.-W. Kim, Y.-J. Lee *et al.*, “Design and creation of dysarthric speech database for development of QoLT software technology,” in *Oriental COCOSDA*. IEEE, 2011, pp. 47–50.
 - [24] F. Rudzicz, A. K. Namasivayam, and T. Wolff, “The TORGO database of acoustic and articulatory speech from speakers with dysarthria,” *Language Resources and Evaluation*, vol. 46, no. 4, pp. 523–541, 2012.
 - [25] P. Bell, M. Gales, P. Lanchantin *et al.*, “Transcription of multi-genre media archives using out-of-domain data,” in *2012 IEEE SLT Workshop*. IEEE, 2012, pp. 324–329.
 - [26] N. Jaitly and G. E. Hinton, “Vocal tract length perturbation (VTLP) improves speech recognition,” in *Proc. ICML Workshop on Deep Learning for Audio, Speech and Language*, vol. 117, 2013.
 - [27] N. Kanda, R. Takeda, and Y. Obuchi, “Elastic spectral distortion for low resource speech recognition with deep neural networks,” in *2013 IEEE ASRU Workshop*. IEEE, 2013, pp. 309–314.
 - [28] X. Cui, V. Goel, and B. Kingsbury, “Data augmentation for deep neural network acoustic modeling,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 9, pp. 1469–1477, 2015.
 - [29] T. Ko, V. Peddinti, D. Povey *et al.*, “Audio augmentation for speech recognition,” in *INTERSPEECH*, 2015, pp. 3586–3589.
 - [30] T. Hayashi, S. Watanabe, Y. Zhang *et al.*, “Back-translation-style data augmentation for end-to-end ASR,” in *2018 IEEE SLT Workshop*. IEEE, 2018, pp. 426–433.
 - [31] B. Vachhani, C. Bhat, and S. K. Kopparapu, “Data Augmentation Using Healthy Speech for Dysarthric Speech Recognition,” in *INTERSPEECH*, 2018, pp. 471–475.
 - [32] Y. Jiao, M. Tu, V. Berisha *et al.*, “Simulating dysarthric speech for training data augmentation in clinical speech applications,” in *ICASSP*. IEEE, 2018, pp. 6009–6013.
 - [33] F. Xiong, J. Barker, and H. Christensen, “Phonetic analysis of dysarthric speech tempo and applications to robust personalised dysarthric speech recognition,” in *ICASSP*. IEEE, 2019, pp. 5836–5840.
 - [34] W. Verhelst and M. Roelands, “An overlap-add technique based on waveform similarity (WSOLA) for high quality time-scale modification of speech,” in *ICASSP*. IEEE, 1993, pp. 554–557.
 - [35] D. Povey, A. Ghoshal, G. Boulianne *et al.*, “The Kaldi speech recognition toolkit,” in *2011 IEEE ASRU Workshop*. IEEE, 2011.
 - [36] P. Swietojanski, J. Li, and S. Renals, “Learning hidden unit contributions for unsupervised acoustic model adaptation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 8, pp. 1450–1463, 2016.
 - [37] H. Wakita, “Direct estimation of the vocal tract shape by inverse filtering of acoustic speech waveforms,” *IEEE Transactions on Audio and Electroacoustics*, vol. 21, no. 5, pp. 417–427, 1973.
 - [38] L. Lee and R. C. Rose, “Speaker normalization using efficient frequency warping procedures,” in *ICASSP*. IEEE, 1996, pp. 353–356.
 - [39] E. Eide and H. Gish, “A parametric approach to vocal tract length normalization,” in *ICASSP*. IEEE, 1996, pp. 346–348.
 - [40] J. Driedger and M. Müller, “A review of time-scale modification of music signals,” *Applied Sciences*, vol. 6, no. 2, p. 57, 2016.
 - [41] J. Allen, “Short term spectral analysis, synthesis, and modification by discrete Fourier transform,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 25, no. 3, pp. 235–238, 1977.
 - [42] J. Andén and S. Mallat, “Deep scattering spectrum,” *IEEE Transactions on Signal Processing*, vol. 62, no. 16, pp. 4114–4128, 2014.
 - [43] S. Young, G. Evermann, M. Gales *et al.*, “The HTK book,” *Cambridge university engineering department*, vol. 3, 2006.
 - [44] R. Caruana, “Multitask learning,” *Machine learning*, vol. 28, no. 1, pp. 41–75, 1997.
 - [45] *Sox, audio manipulation tool*. (accessed Feb 10, 2020). [Online]. Available: <http://sox.sourceforge.net/>