



Domain Adversarial Neural Networks for Dysarthric Speech Recognition

Dominika Woszczyk¹, Stavros Petridis¹, David Millard²

¹ Department of Computing, Imperial College London, UK

² School of Electronics and Computer Science, University of Southampton, UK

d.woszczyk19@imperial.ac.uk, stavros.petridis04@imperial.ac.uk, dem@ecs.soton.ac.uk

Abstract

Speech recognition systems have improved dramatically over the last few years, however, their performance is significantly degraded for the cases of accented or impaired speech. This work explores domain adversarial neural networks (DANN) for speaker-independent speech recognition on the UAS dataset of dysarthric speech. The classification task on 10 spoken digits is performed using an end-to-end CNN taking raw audio as input. The results are compared to a speaker-adaptive (SA) model as well as speaker-dependent (SD) and multi-task learning models (MTL). The experiments conducted in this paper show that DANN achieves an absolute recognition rate of 74.91% and outperforms the baseline by 12.18%. Additionally, the DANN model achieves comparable results to the SA model's recognition rate of 77.65%. We also observe that when labelled dysarthric speech data is available DANN and MTL perform similarly, but when they are not DANN performs better than MTL.

Index Terms: speech recognition, domain adversarial training neural networks, dysarthric speech, multi-task learning, transfer-learning

1. Introduction

The human voice is a powerful communication tool that makes it possible for us to interact with others but also with technology. Voice interfaces offer an easy and natural way of communication with devices. Unfortunately, speech recognition systems still perform relatively poorly on dysarthric speech.

Dysarthria is a motor speech disorder that affects an individual's respiration, articulation, phonation and prosody, reducing their ability to produce intelligible speech. It is often the result of a neurological pathology (e.g. Parkinson or Alzheimer) but it can also be the consequence of any brain-related injury. The impairment can take a multitude of patterns, each one depending on the severity and the cause of the disorder. However, this condition generally affects the tempo, rhythm or volume of an individual's speech. Due to the numerous causes of the disorder, dysarthria is classified into different types with distinguishable patterns, albeit each person has their differences. Nevertheless, the mistakes are consistent and predictable, while it could be hard to differentiate for a human ear, it makes it possible for a machine to adapt to an individual's speech.

With the advances of deep learning, automatic speech recognition (ASR) has reached high accuracy and robustness but still performs poorly on impaired speech. In a review of [1], the authors discuss the poor performance of ASR systems on dysarthric speech and call for necessary future work in this domain. In [2], they show that current voice assistants reach on average an accuracy of 50-60% on impaired speech while the minimal satisfactory rate for healthy speech is considered to be 90-95% and is considered tolerable between 65% to 80% for

individuals with speech impairment [3]. The main issue is that deep learning architectures, with numerous parameters to optimise, need to train on a lot of data to achieve a good generalisation. Common ASR systems are trained on thousands of hours of healthy speech and often overlook speech with particularities such as accents or impairments. Although many databases exist for dysarthric speech for multiple languages such as English ([4, 5, 6]), Cantonese ([7]) or Dutch ([8]), they have a limited number of speakers and/or utterances. This is due to the difficulty of collecting data from individuals as they can be difficult to access and the effort required to produce large enough vocabulary is more exhausting than for healthy individuals. Looking into this problem, different domain adaptation architectures have been studied and successfully applied to dysarthric speech such as speaker-dependent (SD) and speaker-adaptive (SA) models. SA models showed particularly promising results [9, 10, 11]. Nevertheless, those approaches still depend on available labelled data of the target speaker. Adversarial training using domain adversarial training of neural network (DANN) has been successfully applied to unsupervised domain adaptation in computer vision [12] and ASR [13, 14], making it an interesting approach to apply to dysarthric speech domain adaptation.

This work investigates the DANN model on the task of 10 digits classification on the Universal Access Speech (UAS) dysarthric speech dataset. The model is compared to the performance of a multi-task learning model (MTL) as well as speaker-dependent and speaker-adaptive systems on an end-to-end baseline. The main contributions of the paper are as follows:

- To the authors' knowledge, this is the first work to apply end-to-end DANN architecture on the problem of dysarthric speech recognition and to compare it to other known solutions. We show that the model yields an overall word recognition rate (WRR) improvement of 12.18% over the baseline.
- We evaluate the results in the light of the satisfactory rate for dysarthric speech recognition (65%) and observe that it is reached for 11 out of 15 speakers.

2. Related work

2.1. Dysarthric speech

Traditional machine learning is characterised by training and testing on data having the same distribution. However, collecting enough data is often either too hard or too expensive. This creates a difference of *domains* between the data used. In the case of an ASR system, the discrepancy can be caused by different pronunciations, and/or background noise. Multiple methods have been studied for dysarthric speech adaptation. One method is data augmentation with mono or multi-lingual healthy speech [15] or simulated dysarthric speech by altering the tempo and the pitch [16]. Alternatively, dysarthric

speech can be turned into healthy speech using Generative Adversarial Networks (GANs) [17]. In [18], Mel Frequency Cepstral Coefficients (MFCCs) are used with an artificial neural network (ANN) model to classify digits and silence from the UAS dataset. They achieve an overall WRR of 68.38% for their speaker-independent model on three selected speakers (M07, M05, F05). In later work [19], they extend the model to a multi-view multi-net ANN which is an ensemble model of several ANNs and reach an overall WRR of 85%. Other work also investigated the performance of speaker-adaptive models [9, 10]. Recently, the Euphonia project led by Google I/O collected through their platform a new database of non-standard speech to improve speech recognition for accented and/or impaired speech [20]. In their paper [11], they test sequence-to-sequence architectures, Bidirectional RNN Transducer (RNN-T) and Listen-Attention-Spell (LAS), trained on normal speech and fine-tuned on non-standard speech. They show that the adapted models perform better for both accented and dysarthric speech on all architectures. Additionally, the selection of the source data is important as it will affect the performance depending on the speakers' pronunciation patterns. The authors in [21] propose a model that is selected based on the speaker severity level which significantly improved the adaptation performance over the baseline model. In [22], they used the entropy of posterior probability to select the data to train the model on before the adaptation step.

2.2. Domain adversarial training

Domain adaptation models attempt to extract domain-invariant features that would generalise well on unseen domains, and minimise the distance between source and target domains. There has been many studies on DNN domain adaptation such as work based on factorised hidden layer [23, 24], i-vectors [25], KL-divergence regularisation [26], or knowledge distillation models [27, 28]. Domain adversarial training neural networks (DANN), introduced in [29, 12], is an unsupervised domain adaptation approach that tries to extract domain invariant features such that the model cannot determine whether the data belong to a domain or the other. They propose a reversal gradient layer (GRL) which introduces a negative gradient that maximises the loss of the classifier. The GRL can be easily implemented to any feed-forward network and without any additional parameter. This model provides the best improvements in supervised and semi-supervised setting and is similar to multi-task learning for supervised-learning. A multi-task learning (MTL) model is trained on multiple tasks simultaneously. With the introduction of an auxiliary task, the model is optimised to find parameters that perform well on more than one task, thus leading to better generalisation. In [30], the paper explores the multi-task model and accent embeddings to increase the robustness of ASR on unseen accents. In that model, the secondary task is an accent classifier while the main task is a multi accent acoustic model. They demonstrated that the MTL improves the performances over the single-stream network. Additionally, [13] adapted the DANN model to accented speech and compared it to results given by an MTL architecture. They conclude that the DANN model outperforms the MTL model significantly when trained on unlabelled accented data. We build upon previous work by exploring DANN and MTL for speaker-independent dysarthric speech recognition and compare it to previously established models and scenarios. Additionally, instead of using MFCCs or spectrograms as input features, the features from raw audio files are directly extracted by 1D-CNN

layers as part of an end-to-end system.

3. Methodology

3.1. Dataset

The dysarthric speech was retrieved from the Universal Access Speech (UAS) database [5], due to the consistency, accessibility and the higher number of speakers than other datasets. The dataset contains a total of 765 isolated words for each one of the 15 speakers (3 female and 12 male) with cerebral palsy and with different levels of severity: high, moderate and mild. Each speaker has either spastic¹, athetoid² or mixed type of cerebral palsy. Each word in the vocabulary is also repeated by 13 control speakers (4 female and 9 male). The collected data contains three batches of seven repetitions of digits, common words, 300 distinct uncommon words, computer commands and radio alphabet. For each speaker, their intelligibility level (0-100) has been computed by human listeners. To simplify the problem and training complexity, we retain the ten spoken digits only.

3.2. Models

3.2.1. Baseline model

Convolutional neural networks (CNN) are often used to extract features and have been successfully applied in computer vision for pattern recognition, object detection or semantic segmentation. Additionally, CNN can also be applied to audio data by taking the signal in its raw form, or as a spectrogram or MFCCs. Work by [31] and [32] show that raw features can perform similarly to the pre-processed forms. RNN - based models have been successful on ASR tasks [11]. However, the disadvantage of recurrent architectures is that they are more complex to train and can overfit on smaller datasets. We chose our baseline model to be a simpler 1D-CNN with seven layers and takes raw audio data as input, taking the whole word at once. The resulting features are fed to a fully connected layer with a sigmoid as an activation function. We selected the layers and parameters through experimentation (using grid search) on control and dysarthric speech independently, optimising for the WRR.

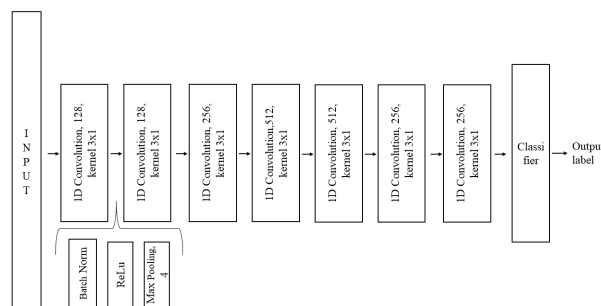


Figure 1: CNN baseline architecture

3.2.2. Domain adversarial training & multi-task learning

Following the work by [12], we implement a gradient reversal layer (GRL) which outputs a negative gradient that al-

¹Stiffness of the muscles which makes movement difficult. The most common type of cerebral palsy.

²Movement disorder caused by damage to the developing brain.

lowers the maximisation of the domain classification loss. The DANN/MTL model we implemented has a shared block of seven convolutional layers just like the baseline as shown in Figure 2. The model then branches out, one of the branches being the label classifier for the actual task trained on both the source and target domain. The second branch is the domain classifier with the GRL followed by a fully connected layer for the domain classification.

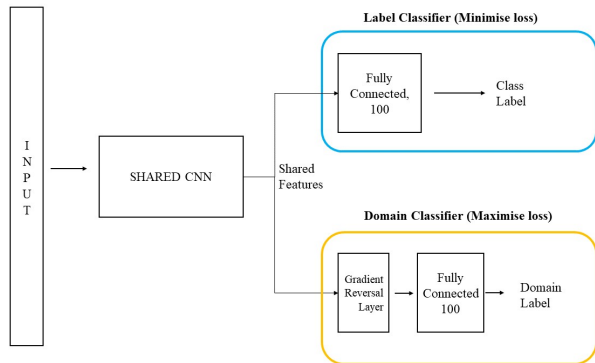


Figure 2: Domain adversarial neural network architecture. The label classifier outputs 0-9 digits. The domain classifier returns 0 for healthy speech and 1 for dysarthric speech.

The objective function of the task is given in (1), where $N = n + n'$ is the total number of input samples and n, n' represent the control samples and dysarthric samples, $\theta_f, \theta_d, \theta_y$ are the parameters of the feature extractors, the domain classifier and the main task classifier. The parameters y represents the classification label, d domain labels and the indexes s, t represent source (healthy) and target (dysarthric) domain respectively. The parameter α is a boolean such that for $\alpha = 1$ we train the label classifier on target domain data as well.

$$E(\theta_f, \theta_y, \theta_d) = \frac{1}{n} \sum_{s=1}^n \mathcal{L}_y^s(\theta_f, \theta_y) + \alpha \left(\frac{1}{n'} \sum_{s=n+1}^N \mathcal{L}_y^t(\theta_f, \theta_y) \right) - \lambda \left(\frac{1}{n} \sum_{s=1}^n \mathcal{L}_d^s(\theta_f, \theta_d) + \frac{1}{n'} \sum_{t=n+1}^N \mathcal{L}_d^t(\theta_f, \theta_d) \right) \quad (1)$$

The MTL model follows the same architecture, with a normal layer instead of the GRL by setting $\lambda < 0$. Through experimentation, we set λ to 1.5 and -0.5 for the DANN and MTL models.

Dysarthric speakers having their own particular speech patterns, each one of them could be a domain of their own. However, in this work, we consider all dysarthric speakers as one domain.

3.3. Evaluation metric

The word recognition rate (WRR) is used as the evaluation metric and is given by (2) as

$$WRR = \frac{\text{word correctly recognised}}{\text{words attempted}} * 100 \quad (2)$$

4. Experiments & results

For all experiments, the same seed was used for the sake of reproducibility. The models were run for 30 epochs with early stopping performed using a hold-out validation set. We set the

batch size to 32. For each speaker, the experiments were repeated three times rotating the batches of iterations and the average WRR was recorded.

4.1. Training with unlabelled dysarthric speech

In the first experiment, we want to see how DANN, MTL and the baseline perform when we do not have any labelled data for dysarthric speech available. We train the baseline on the control speech only (2730 samples) and test on all the repetitions for each dysarthric speaker (210 samples per speaker). For the MTL and DANN models, we set $\alpha = 0$ and we train them on control speech and all dysarthric speakers' unlabelled iterations excluding the test subject, for which the first batch of iterations (70 samples) was used as a validation set and the two other batches were used in the test set (140 samples). The results on Table 1 show that the DANN model performs better than the baseline and the MTL model on average for all groups. One can see on Figure 3 that both DANN and MTL have a lesser variability than the baseline as well as DANN have a higher median than both the baseline and the MTL model.

Table 1: Recognition rates for the baseline, MTL and DANN on dysarthric speakers when trained with labeled control speech and unlabelled dysarthric speech. The number in bold is the best WRR for its row.

Severity	Speaker	Baseline	DANN	MTL
Mild	F05	87.62	85.71	88.10
	M08	96.67	97.14	80.48
	M10	96.19	98.10	88.57
	M14	92.86	88.10	88.57
	M09	25.71	75.24	55.24
Moderate	M11	29.44	53.89	54.44
	F04	38.10	56.67	51.90
	M05	37.62	77.62	61.90
High	M16	48.33	76.67	72.22
	F02	40.48	45.24	41.90
	M07	25.24	61.43	46.67
	M01	22.73	48.18	51.82
	M12	21.11	28.89	23.33
	F03	14.29	31.90	34.76
	M04	29.53	39.60	42.28
	Mean (SD)	Mild	79.81 (30.46)	88.86 (9.35)
	Mod	35.05 (4.87)	62.73 (12.97)	56.08 (5.20)
	High	28.82 (11.82)	47.42 (16.85)	44.71 (15.17)
	All	47.06 (30.14)	64.29 (23.00)	58.81 (20.63)




Figure 3: WRR distributions for the baseline, MTL and DANN on dysarthric speakers when trained with labelled control speech and unlabelled dysarthric speech.

4.2. Training with labelled dysarthric speech

The second set of experiments are conducted to compare the performance of the DANN, MTL and the baseline model as a speaker-independent (SI) model (none of the utterances of the speaker was present in the training set) when some labelled dysarthric speech is available. In this case, we set $\alpha = 1$ for

both DANN and MTL. The baseline model was trained on all control speakers (2730 samples) and the first batch of iterations of the tested dysarthric speaker was used as a validation (70 samples) set and the two other batches were used in the test set (140 samples).

Additionally, we compare the models to the baseline when it is trained as a speaker-dependent (SD) and speaker-adaptive (SA) system. The SA model is the SI model fine-tuned with one of the three iterations of the tested speaker and validated and tested with the two remaining iterations. Through experimentation, fine-tuning only the last convolutional layer and the linear classifier yields the best results. Therefore those are the ones shown in Table 2. The SD model was trained on control speakers combined with one iteration of the tested subject (2800 samples), the other two iterations were used for validation (70 samples) and testing (70 samples).

Table 2: Recognition rates for models trained on labelled dysarthric and control speech. The number in bold is the best WRR for its row.

Severity	Speaker	SI	DANN	MTL	SD	SA
Mild	F05	86.67	90.48	90.48	89.21	92.86
	M08	95.24	95.24	98.10	96.19	94.29
	M10	100.00	98.58	98.00	98.86	100.00
	M14	100.00	94.29	94.00	96.10	100.00
	M09	53.57	92.14	90.00	78.57	94.29
Moderate	M11	50.00	88.89	85.56	74.82	85.56
	F04	67.14	75.71	77.81	73.55	72.86
	M05	70.48	86.19	89.52	82.06	100.00
High	M16	65.56	76.67	82.22	74.82	80.00
	F02	53.33	48.10	45.71	49.05	72.86
	M07	45.24	75.24	72.38	64.29	84.29
	M01	41.82	62.73	65.45	56.67	53.33
	M12	30.56	25.00	26.11	27.22	45.00
	F03	50.48	70.00	71.43	63.97	50.00
	M04	30.87	44.30	44.97	40.05	40.00
	Mean (SD)					
Mild	87.10 (23.11)	94.15 (3.50)	94.12 (4.94)	91.79 (10.09)	96.29 (5.35)	
Mod	62.54 (10.99)	83.60 (6.96)	84.30 (5.96)	76.81 (4.59)	85.95 (13.58)	
High	45.41 (12.51)	57.43 (19.08)	58.32 (19.86)	53.72 (16.24)	60.78 (17.89)	
All	62.73 (23.56)	74.91 (21.60)	75.45 (21.59)	71.03 (21.12)	77.65 (21.18)	




Figure 4: WRR distributions for all severity subjects for models trained with labelled control and dysarthric speech.

The results on Table 2 show that the baseline SI model reached an overall WRR of $(62.73 \pm 23.56)\%$ and is the worst-performing model for almost all speakers, with some exceptions. Due to the small number of training samples, the model overfits relatively fast. Each dysarthric speaker has a unique impairment which makes it harder for a model to generalise well on multiple speakers when it has not been introduced to the test subject specifically. However, this model achieved similar mean WRR of 67.46% for the M05, F05 and M07 speakers as the SI model from [18] (68.38%), which used carefully selected MFCCs.

On the other hand, the best performing system is the SA model, reaching the highest WRR reaching an average of $(77.65 \pm 21.18)\%$, and severity specific means of 96.28%, 85.96% and 60.78% for mild, moderate and high levels. This confirms pre-

vious work [9, 10, 11] that showed that by fine-tuning on a target subject it is possible to improve greatly the performance of the ASR system. However, Figure 4 shows that there is a high variability in the performance comparable to the SI model and greater than for the other models. Table 2 also shows that both DANN and MTL models outperform the baseline model for both SD and SI scenarios. Moreover, both models reach similar performance to the speaker adaptive system for all severity levels, outperforming it at times. However, the DANN and MTL models perform similarly as shown in Figure 4 but the MTL model has a higher median and a lower variance indicating a more consistent performance.

Although MTL has a slightly better average WRR (0.54% difference), none of the models is significantly better than the other. Both models introduce regularisation by extraction domain invariant features, with the difference being that the DANN model attempts to confuse the domain classifier and maximise the loss while MTL is minimising its loss. In this scenario, MTL achieves a slightly higher mean. In terms of the SD model which achieved on average a WRR of $(71.03 \pm 21.12)\%$, it performs better than the SI one but is outperformed by the SA model. Because of the small training set, the SD system was trained on, it is easier to overfit. However, the SA model benefits from the pre-training on the bigger dataset. Similarly, the SD model is outperformed by the MTL and DANN models. Those models also benefit from more generalised features as they were trained in a multi-task manner and on a bigger dataset as well.

Regarding usability, from results on Table 2 we can see that the SA, DANN and MTL models achieve satisfactory recognition rates at healthy speech level (90%) for all speakers with mild severity dysarthria. The systems also reach the tolerable level for dysarthric users for all subjects at moderate severity with no WRR below 70%. For high severity speakers, all reach three out of seven. In total, all three best-performing algorithms achieve a satisfactory WRR for 11 out of 15 speakers.

5. Conclusion

Speech recognition has made great progress with the advances of deep learning architectures. However, it still lacks robustness to non-standards speech for which there is usually less labelled data available. This paper leverages domain-invariant features with domain adversarial training to cope with limited data on the UAS dataset of dysarthric speech. In the conducted experiments, domain adversarial training neural networks (DANN) and multi-task-learning (MTL) models were studied for cases when some labelled dysarthric data is and is not available. We compare them to a baseline CNN model trained on speaker-dependent (SD) and speaker-independent (SI) scenarios and a speaker-adaptive (SA) system by fine-tuning each model on each speaker. Results show that both the DANN architecture can significantly improve over the baseline CNN for speakers of all levels of severity when some labelled data is available. Moreover, the results are similar to the best performing model fine-tuned on each speaker (SA), outperforming it for some subjects. We also observe that in this supervised scenario MTL and DANN perform very similarly but the DANN model outperforms the MTL model when trained with no target domain data. Furthermore, our end-to-end baseline achieves similar WRR to a model from previous work using MFCCs. Finally, we discuss the usability of the models and observe that the DANN model achieve a satisfactory WRR for dysarthric speakers for 11 out of 15 subjects.

6. References

- [1] M. Moore, H. Venkateswara, and S. Panchanathan, "Whistle-blowing asrs: evaluating the need for more inclusive automatic speech recognition systems," in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 2018, 2018, pp. 466–470.
- [2] F. Ballati, F. Corno, and L. De Russis, "'hey siri, do you understand me?': Virtual assistants and dysarthria," in *Intelligent Environments (Workshops)*, 2018, pp. 557–566.
- [3] M. Jefferson, "Usability of automatic speech recognition systems for individuals with speech disorders: Past, present, future, and a proposed model," 2019.
- [4] 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.
- [5] 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.
- [6] X. Menendez-Pidal, J. B. Polikoff, S. M. Peters, J. E. Leonzio, and H. T. Bunnell, "The nemours database of dysarthric speech," in *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP'96*, vol. 3. IEEE, 1996, pp. 1962–1965.
- [7] K. H. Wong, Y. T. Yeung, E. H. Chan, P. C. Wong, G.-A. Levow, and H. Meng, "Development of a cantonese dysarthric speech corpus," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [8] E. Yilmaz, M. Ganzeboom, L. Beijer, C. Cucchiari, and H. Strik, "A dutch dysarthric speech database for individualized speech therapy research," in *Proceedings of the 10th International Conference on Language Resources and Evaluation*. Portoroz, Slovenia: European Language Resources Association, 2016, pp. 792–795.
- [9] F. Chen and A. Kostov, "Optimization of dysarthric speech recognition," in *Proceedings of the 19th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Magnificent Milestones and Emerging Opportunities in Medical Engineering (Cat. No. 97CH36136)*, vol. 4. IEEE, 1997, pp. 1436–1439.
- [10] F. Rudzicz, "Comparing speaker-dependent and speaker-adaptive acoustic models for recognizing dysarthric speech," in *Proceedings of the 9th international ACM SIGACCESS conference on Computers and accessibility*. ACM, 2007, pp. 255–256.
- [11] 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.
- [12] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [13] S. Sun, C.-F. Yeh, M.-Y. Hwang, M. Ostendorf, and L. Xie, "Domain adversarial training for accented speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4854–4858.
- [14] A. Tripathi, A. Mohan, S. Anand, and M. Singh, "Adversarial learning of raw speech features for domain invariant speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5959–5963.
- [15] Y. Takashima, T. Takiguchi, and Y. Ariki, "End-to-end dysarthric speech recognition using multiple databases," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6395–6399.
- [16] B. Vachhani, C. Bhat, and S. K. Kopparapu, "Data augmentation using healthy speech for dysarthric speech recognition," in *Interspeech*, 2018, pp. 471–475.
- [17] S. H. Yang and M. Chung, "Improving dysarthric speech intelligibility using cycle-consistent adversarial training," *arXiv preprint arXiv:2001.04260*, 2020.
- [18] S. R. Shahamiri and S. S. B. Salim, "Artificial neural networks as speech recognisers for dysarthric speech: Identifying the best-performing set of mfcc parameters and studying a speaker-independent approach," *Advanced Engineering Informatics*, vol. 28, no. 1, pp. 102–110, 2014.
- [19] S. R. Shahamiri and S. S. B. Salim, "A multi-views multi-learners approach towards dysarthric speech recognition using multi-nets artificial neural networks," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 5, pp. 1053–1063, 2014.
- [20] "How ai can improve products for people with impaired speech," <https://blog.google/outreach-initiatives/accessibility/impaired-speech-recognition/>, 2019.
- [21] M. J. Kim, J. Yoo, and H. Kim, "Dysarthric speech recognition using dysarthria-severity-dependent and speaker-adaptive models," in *Interspeech*, 2013, pp. 3622–3626.
- [22] 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)*, 2020, pp. 7424–7428.
- [23] L. Samarakoon, B. Mak, and A. Y. S. Lam, "Domain adaptation of end-to-end speech recognition in low-resource settings," in *2018 IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 382–388.
- [24] K. C. Sim, A. Narayanan, A. Misra, A. Tripathi, G. Pundak, T. N. Sainath, P. Haghani, B. Li, and M. Bacchiani, "Domain adaptation using factorized hidden layer for robust automatic speech recognition," in *Interspeech*, 2018, pp. 892–896.
- [25] V. Gupta, P. Kenny, P. Ouellet, and T. Stafylakis, "I-vector-based speaker adaptation of deep neural networks for french broadcast audio transcription," in *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2014, pp. 6334–6338.
- [26] D. Yu, K. Yao, H. Su, G. Li, and F. Seide, "Kl-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 7893–7897.
- [27] M. Orbes-Arteainst, J. Cardoso, L. Sørensen, C. Igel, S. Ourselin, M. Modat, M. Nielsen, and A. Pai, "Knowledge distillation for semi-supervised domain adaptation," in *OR 2.0 Context-Aware Operating Theaters and Machine Learning in Clinical Neuroimaging*. Springer, 2019, pp. 68–76.
- [28] T. Asami, R. Masumura, Y. Yamaguchi, H. Masataki, and Y. Aono, "Domain adaptation of dnn acoustic models using knowledge distillation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5185–5189.
- [29] H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, and M. Marchand, "Domain-adversarial neural networks," *arXiv preprint arXiv:1412.4446*, 2014.
- [30] A. Jain, M. Upreti, and P. Jyothi, "Improved accented speech recognition using accent embeddings and multi-task learning," in *Interspeech*, 2018, pp. 2454–2458.
- [31] R. A. Solovyev, M. Vakhrushev, A. Radionov, V. Aliev, and A. A. Shvets, "Deep learning approaches for understanding simple speech commands," *arXiv preprint arXiv:1810.02364*, 2018.
- [32] W. Dai, C. Dai, S. Qu, J. Li, and S. Das, "Very deep convolutional neural networks for raw waveforms," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 421–425.