



Exploiting Cross-Domain And Cross-Lingual Ultrasound Tongue Imaging Features For Elderly And Dysarthric Speech Recognition

Shujie Hu¹, Xurong Xie², Mengzhe Geng¹, Mingyu Cui¹, Jiajun Deng¹, Guinan Li¹, Tianzi Wang¹, Helen Meng¹, Xunying Liu¹

¹The Chinese University of Hong Kong, Hong Kong SAR, China

²Institute of Software, Chinese Academy of Sciences, China

{sjhu,mzeng,mycui,jjdeng,gnli,twang,hmmeng,xyliu}@se.cuhk.edu.hk, xurong@iscas.ac.cn

Abstract

Articulatory features (AFs) are inherently invariant to acoustic signal distortion. Their practical application to atypical domains such as elderly, disordered speech across languages is limited by data scarcity. This paper presents a cross-domain and cross-lingual Acoustic-to-Articulatory (A2A) inversion approach that utilizes the parallel audio and ultrasound tongue imaging (UTI) data of the 24-hour TaL corpus in A2A model training before being adapted to three datasets: the English DementiaBank Pitt and Cantonese JCCOCC MoCA elderly speech corpora; and the English TORGO dysarthric speech data, to produce UTI based AFs. Experiments suggest incorporating the generated AFs consistently outperforms the baseline TDNN/Conformer ASR systems using acoustic features only by statistically significant word/character error rate reductions up to 4.75%, 2.59% and 2.07% absolute (14.69%, 10.64% and 22.72% relative) after data augmentation, speaker adaptation and cross system multi-pass decoding.

Index Terms: Articulatory Inversion, Elderly and Dysarthric Speech, End-to-End ASR, Domain Adaptation, Cross-lingual

1. Introduction

Despite the rapid progress of automatic speech recognition (ASR) technologies in the past few decades, accurate recognition of elderly and disordered speech remains a challenging task [1–7]. Neurocognitive disorders, such as Alzheimer’s disease (AD), are often found among older adults and manifest themselves in speech and language impairments including weakened neuro-motor control in speech production and imprecise articulation. ASR based assistive technology developments tendering for such users’ needs play a vital role in not only improving their quality of life and social inclusion, but also facilitating large scale automatic speech based early diagnosis of neurocognitive impairment and preventive care [4, 8–11].

Elderly and disordered speech bring challenges on all fronts to current deep learning based ASR technologies predominantly targeting non-aged, healthy adult users. First, a large mismatch between such data and non-aged, healthy adult voice is often observed. Such difference manifests itself across many fronts including articulatory imprecision, decreased volume and clarity, changes in pitch, increased dysfluencies and slower speaking rate. Second, the co-occurring disabilities, mobility or accessibility limitations often found among elderly and disordered speakers lead to the difficulty in collecting large quantities of such data that are essential for current data intensive ASR system development. In addition, sources of variability commonly found in normal speech including accent or gender, when further compounded with those over age and speech pathology severity, create large diversity among speakers [12, 13].

Human speech production involves the coordinated movements of various articulators such as the tongue, lips, teeth and palate. Articulatory movement features are inherently invariant to extrinsic acoustic distortion. They have been successfully applied to normal [14–16] and pathological speech [17–20] recognition. In practice, recording detailed articulatory movements and vocal tract shape normally requires the use of intrusive electromagnetic articulography (EMA) [21] or magnetic resonance imaging (MRI) [22] techniques. Compared to EMA and MRI, ultrasound tongue imaging (UTI) [23–25] is more portable, less non-invasive and costly. UTI uses B-mode diagnostic ultrasound to visualize the tongue surface movements during speech production at a high frame rate [26]. However, there are currently very few publicly available UTI corpora, all of which were designed for the English and of limited size [26–28]. By far the largest Tongue and Lips corpus (TaL) [26] contains 24 hours of parallel ultrasound, video, and audio data collected from 81 native speakers of English. The practical and wider use of UTI based articulatory features in ASR systems for both normal and atypical speech task domains, and across languages, is hindered by the scarcity of such specialist data.

An alternative and general approach to obtain articulatory information is to estimate it from the more accessible acoustic speech signals using data driven artificial neural network based acoustic-to-articulatory (A2A) inversion techniques [29–32]. As the A2A inversion model training only requires a part of training materials containing parallel acoustic-articulatory data, the resulting inversion model can be used to produce articulatory features when only audio recordings are available. A wider and more practical application of articulatory features in ASR systems becomes possible. Prior researches on A2A inversion were conducted predominantly on normal speech task domains [29–31] and most of them used EMA based articulatory features. In contrast, very limited researches were carried out on A2A inversion for UTI based articulatory features [33].

In order to address the issues mentioned above, this paper presents a cross-domain and cross-lingual A2A inversion approach that utilizes the parallel acoustic-articulatory data of the 24-hour TaL corpus [26] in A2A model training before being cross-domain and cross-lingual adapted to three datasets across two languages: the English DementiaBank Pitt [34] and the Cantonese JCCOCC MoCA [35] elderly speech corpora; and the English TORGO [36] dysarthric speech data to produce UTI articulatory features. Bi-LSTM based deep neural network A2A inversion models are used. A cross-domain adaptation network is used to reduce the acoustic mismatch between the TaL corpus [26] and the target elderly and disordered speech corpora.

The main contributions are summarized below. To the best of our knowledge, this work is the first systematic study of using audio and ultrasound tongue imaging parallel recordings trained

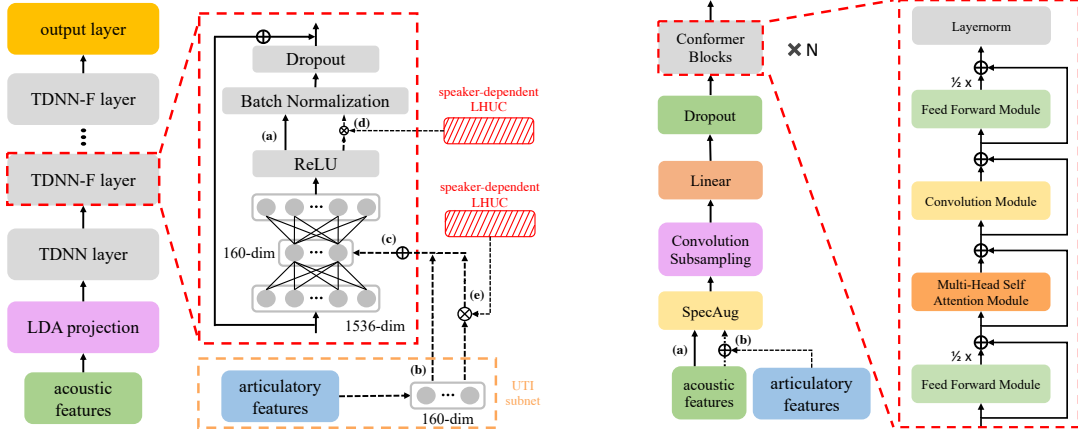


Figure 1: The left part of the Figure is the factored TDNN based ASR and AASR system architecture. Different connection combinations in this figure form different systems. Retaining connection (a) while discarding others leads to the ASR baseline system; Retaining (a), (b) and (c) (for articulatory modality fusion) produces the TDNN AASR system. The connections (d) and (d)+(e) are used for LHUC speaker adaptive training of ASR and AASR systems. The right part of the Figure shows the Conformer based end-to-end system. The AASR system uses the acoustic-articulatory concatenation via connection (b).

A2A inversion models for UTI articulatory features generation targeting elderly and disordered speech recognition. In contrast, prior researches on using UTI features were conducted mainly in the context of non-aged adult or child speech recognition [26, 37–39]. In addition, this work is also the first use of cross-lingual and cross-domain UTI based articulatory features to develop elderly and disordered speech recognition systems. On all three tasks across both the English and Cantonese languages, significant WER or CER reduction is obtained over the acoustic features only baseline hybrid TDNN [40] and E2E Conformer [41] systems after speed perturbation plus SpecAugment based data augmentation, learning hidden unit contributions (LHUC) [42] speaker adaptation and cross system multipass decoding [43] are applied. Further analysis of the RMSE metric and the phonetic discrimination brought by UTI articulatory features via T-SNE [44] visualization are also given.

2. Articulatory Feature Based Speech Recognition

This section describes the time delay neural network (TDNN) [40] or E2E Conformer [41] based ASR and acoustic-articulatory feature based speech recognition (AASR) system architecture on elderly and disordered corpora.

Both of the baseline hybrid TDNN ASR and AASR systems share the same main structure based on a 7-layer or 14-layer factorized TDNN (TDNN-F) model with a semi-orthogonal constraint and they are trained using the sequence discriminative lattice-free MMI (LF-MMI) criterion (the left part of Fig. 1). The TDNN-F hidden layers positioned after LDA projection and TDNN layer are shown in the red dotted box in Figure 1. The AASR system is produced by adding an ultrasound tongue imaging (UTI) subnet using connections (b) and (c). The output of the UTI subnet is then added to the 160-dimensional factorized hidden layer. To model the large variability among elderly and disordered speakers, learning hidden unit contributions (LHUC) [42] based speaker adaptive training (SAT) is used. Speaker-level LHUC scaling factors (in red) are applied to the main TDNN-F architecture via connection (d) and UTI component via connection (e) respectively (left part, Fig. 1).

E2E Conformer ASR and AASR system architectures are shown in the right part of Fig. 1, where the details of Conformer blocks are shown in the red dotted box. UTI articulatory

features are incorporated at the input layer by connection (b).

3. Acoustic-To-Articulatory Inversion

3.1. In-domain A2A Inversion

The objective of acoustic-to-articulatory inversion is to learn the mapping between acoustic and UTI based articulatory feature representations. To this end, a suitable inversion model is required. In previous research, state-of-the-art A2A performance was achieved using long short term memory (LSTM) networks [31, 45]. In this paper, the out-of-domain non-aged normal speech of the TaL dataset [26] accompanied with parallel ultrasound data collected from 81 native speakers of English is used to construct Bi-LSTM based A2A inversion model. The efficacy of the A2A inversion model is evaluated by measuring the performance of the resulting AASR systems built using either the ground truth, original UTI features¹, or those generated using A2A inversion. The UTI based articulatory features are used in both AASR system training and evaluation stages.

Table 1: WER (%) of ASR and AASR systems on the TaL test set.

ASR	WER (%)	
	AASR	
	original features	generated features
8.19	7.47	7.48

The results in Table 1 illustrate that: 1) Using either the original or inverted articulatory features produce performance improvements over the ASR system; 2) the AASR system using the inverted articulatory features has comparable performance compared to that using original articulatory features.

3.2. Cross-Domain And Cross-Lingual A2A Inversion

Due to the acoustic domain and language mismatch, a direct cross-domain and cross-lingual application of the A2A inversion model trained on a normal, non-aged acoustic-articulatory parallel data to the elderly and disordered speech is problematic, as was shown in the previous researches on cross-domain audio-visual inversion [3] and A2A inversion [46]. To this end, the large acoustic domain and language mismatch are minimized using multi-level adaptive networks (MLAN) [3, 46, 47] before A2A inversion can be performed.

¹The UTI based articulatory features are extracted following [26] where the 2D Discrete Cosine Transform (DCT) with 12×12 coefficients to each resized ultrasound frame was applied.

The MLAN training process includes the following steps: 1) the first level DNN is trained with the data from the in-domain English or Cantonese language specific elderly or disordered speech; 2) the resulting in-domain speech trained DNN is then used to produce bottleneck features for the out-of-domain data of the English TaL dataset; 3) the second level DNN is trained using the out-of-domain TaL audio data concatenated with the bottleneck features computed from the previous step. The combined effect produced by these two cascaded component DNNs is such that, when feedforwarding the elderly or dysarthric data of a target language into the resulting MLAN network, the final bottleneck features produced at the second-level DNN component will exhibit smaller mismatch against those obtained by feedforwarding the English TaL acoustic data into the MLAN network. The resulting cross-domain and cross-lingual adapted bottleneck features are used in A2A inversion model training and articulatory feature generation for elderly or disordered audio data of English or Cantonese. Further detailed descriptions of the cross-domain MLAN model are in [3,46,47].

4. Experiments

4.1. Experiments On Elderly Speech

1) *The DementiaBank Pitt corpus*: The DementiaBank Pitt [34] corpus contains about 33-hour audio recorded over interviews between the 292 elderly participants and clinical investigators. After silence stripping [1], the training set contains 15.7 hours of audio data (29682 utterances) while the development and evaluation sets contain 2.5 hours (5103 utterances) and 0.6 hours (928 utterances) of audio respectively. After a combination of speaker independent of elderly speech and dependent speed perturbation [48] of non-aged investigators’ speech based data augmentation, the duration of training data is increased to 58.9 hours (112830 utterances).

2) *The JCCOCC MoCA Corpus*: The Cantonese JCCOCC MoCA corpus [35] is based on speech recordings of cognitive impairment assessment interviews between 256 elderly participants and clinical investigators. After removal of excessive silence, the training set contains 32.1 hours of speech (95448 utterances) while the development and evaluation sets contain 3.5 hours (13675 utterances) and 3.4 hours (13414 utterances) of speech respectively. Data augmentation similar to those adopted on the DementiaBank Pitt corpus produces a 156.9 hours augmented training set (389409 utterances).

3) *Experiment Setup for the DementiaBank Pitt Corpus and the JCCOCC MoCA Corpus*: In our experiments, the 14-layer LF-MMI based TDNN-F models (17M) as shown in Figure 1 are implemented using the Kaldi toolkit [49] while Conformer based end-to-end systems² (52M) are implemented using the Espnet toolkit [50]. A 3-frame context window is used in both ASR and AASR hybrid LF-MMI trained TDNN systems. 40-dimensional Mel-scale filter banks (FBKs) are used as the input acoustic features. The 144-dimensional A2A generated UTI based articulatory features are produced by the approach described in section 3.2. For both the hybrid phonetic TDNN and E2E graphemic (character) Conformer systems, a word level 4-gram LM is used in recognition.

4) *Performance Analysis*: Domain adaptation is essential in transferring articulatory features from non-aged TaL speech to elderly corpora (Sys. 3 vs. 2). After incorporating the cross domain and cross lingual inverted UTI based articulatory features obtained on the DementiaBank Pitt or the JCCOCC MoCA au-

²12 encoder+12 decoder layers, feed-forward layer dim=2048, attention heads=4, attention heads dim=256, CTC+AED cost.

dio data, the resulting AASR systems shown in Table 2 consistently outperform the comparable baseline ASR systems using acoustic features only by a significant margin in WER or CER (Sys. 3 vs. 1; Sys. 5 vs. 4; Sys. 7 vs. 6; Sys. 10 vs. 9) before and after data augmentation and LHUC speaker adaptation are applied. The performance improvements obtained by incorporating the inverted UTI features are retained after increasing the model complexity of the baseline TDNN ASR systems to be comparable to that of the AASR systems³ (e.g. Sys. 4.L vs. 5). In particular, significant WER reductions of 2.39%-2.43% are obtained on the elderly participants of the development and evaluation sets of the DementiaBank Pitt after data augmentation and LHUC adaptation of TDNNs (Sys. 7 vs. 6). Similar consistent improvement of 0.56%-1.36% on the JCCOCC MoCA corpus can be obtained (Sys. 7 vs. 6).

Further investigation on the alternative forms of acoustic-articulatory modality fusion suggests that for both TDNN and Conformer based AASR systems, incorporating UTI based articulatory features via “score A+AA”, a linear interpolation of system specific log-likelihood scores between ASR and AASR systems (Sys. 6 + 7, Sys. 9 + 10), produces further performance improvements over hidden layer or input layer feature fusion (Sys. 8 vs. 7, 11 vs. 10). After cross system 2-pass decoding [43], which uses the Conformer system (Sys. 10) to rescore the score fused AASR TDNN system (Sys. 8) outputs, overall WER and CER reductions of 4.75% and 2.59% over the baseline TDNN systems are obtained (Sys. 12 vs. 6) on the DementiaBank Pitt and the JCCOCC MoCA corpora respectively.

Further ablation studies suggest that the best TDNN-F hidden layer to fuse the cross-domain and cross-lingual inverted UTI features vary across the English DementiaBank Pitt and the Cantonese JCCOCC MoCA data (7th vs. 3rd hidden layer). This may be attributed to the fact that the lower positioned TDNN-F layers model acoustic variability while the higher ones are more related to language specific phonetic targets.

4.2. Experiments On Disordered Speech

1) *the TORGO Corpus*: The TORGO dataset is a disordered speech corpus with 8 dysarthric and 7 control speakers. All 7 control speakers’ data together with two thirds of the 8 dysarthric speakers’ data are merged into the training set while the remaining dysarthric speech serves as the test data. After silence stripping and a combination of disordered speaker independent and control speaker dependent speed perturbation [48] based data augmentation, the total training set contains 34.11 hours of data while the test set has 1.02 hours of speech.

2) *Experiment Setup for the TORGO Corpus*: The hybrid LF-MMI trained TDNN-F ASR system contains 7 hidden layers (9M), while the Conformer models⁴ (18M) are trained to directly model grapheme (character) sequence outputs. A 3-frame context window is used in both ASR and AASR hybrid TDNN-F systems. 40-dimensional Mel-scale filter banks (FBKs) and 144-dimensional UTI based features are used as the input acoustic and articulatory features respectively. A 3-gram LM trained by all the TORGO transcripts is used in decoding.

3) *Performance Analysis*: Several trends similar to those found on the two elderly speech tasks can be observed. Compared with ASR systems, a significant overall WER reduction of 1.02% absolute (11.20% relative) can be obtained after data

³#paramters of TDNN ASR system (Sys. 4) and larger one (Sys. 4.L) are 17M and 23M respectively, while that of TDNN AASR system (Sys. 5) is 23M, where the size of MLAN and inversion model is 6M.

⁴8 encoder+4 decoder layers, feed-forward layer dim=1024, attention heads=4, attention heads dim=256, CTC+AED cost.

Table 2: The performance of baseline ASR and AASR systems using the cross-domain/lingual inverted UTI articulatory features (AFs) on 1) the English DementiaBank Pitt development (Dev) and evaluation (Eval) sets; 2) the Cantonese JCCOCC MoCA development (Dev) and evaluation (Eval) sets containing elderly speakers only; AND 3) the English dysarthric TORGO test set. “INV” and “PAR” refer to clinical investigator and elderly participant. “Seve.,” “Mod.” and “Mild” are intelligibility levels of the dysarthric speakers: severe, moderate and mild. “4.L” denotes a larger TDNN model with comparable #parameters as the AASR system (Sys. 5). AFs incorporated using “hidden layer” fusion and optionally further score fusion “score A+AA”. “→” stands for multi-pass decoding. † denotes statistical significant (MAPSSWE, $\alpha = 0.05$) differences obtained against the baseline ASR systems (Sys. 1, 4, 6, 9)

Sys.	Model	MLAN	AA fusion	Data Aug.	LHUC	WER (%)				CER (%)			WER (%)				
						DementiaBank Pitt				Overall	JCCOCC MoCA		TORGO				
						Dev		Eval			Dev	Eval	Overall	Seve.	Mod.	Mild	Overall
PAR.	INV.	PAR.	INV.	PAR.	INV.	PAR.	INV.	Seve.	Mod.	Mild	Overall						
1		✗	✗			50.71	21.58	39.16	21.20	36.04	30.89	27.95	29.41	16.22	10.31	3.87	11.62
2		✗	hidden layer	✗	✗	51.82	20.70	39.87	20.87	36.24	30.88	28.27	29.57	16.67	10.00	3.64	11.73
3		✓	hidden layer			48.83[†]	19.87[†]	37.31[†]	20.64	34.28[†]	28.86[†]	26.15[†]	27.50[†]	16.14	6.12	3.48	10.61[†]
4	TDNN	✗	✗			47.93	19.91	36.66	19.76	33.80	26.87	23.71	25.28	12.80	8.78	3.64	9.47
4.L		✗	✗	✓	✗	47.80	21.31	36.85	20.53	34.37	27.02	24.06	25.53	13.74	9.39	2.94	9.97
5		✓	hidden layer			45.82[†]	19.21[†]	34.89[†]	18.42	32.35[†]	25.06[†]	22.88[†]	23.96[†]	12.80	4.59	2.71	8.35[†]
6		✓	✗			45.49	19.26	35.44	18.42	32.33	25.77	22.94	24.35	12.52	8.27	3.25	9.11
7		✓	hidden layer	✓	✓	43.06[†]	18.31[†]	33.05[†]	18.09	30.57[†]	24.41[†]	22.38	23.39[†]	12.20	5.00	2.63	8.09[†]
8		✓	score A+AA			41.97[†]	17.76[†]	31.83[†]	17.43	29.69[†]	23.69[†]	21.18[†]	22.43[†]	11.79	5.10	2.63	7.90[†]
9		✗	✗			48.71	20.97	36.93	19.42	34.57	33.08	31.24	32.15	21.22	6.63	4.80	13.72
10	Conformer	✓	input layer	✓	✗	47.86	20.54	36.57	17.76	33.95[†]	31.88[†]	30.20	31.04[†]	19.47	6.73	3.72	12.53[†]
11		✓	score A+AA			47.68[†]	20.38[†]	36.13	17.87	33.75[†]	31.78[†]	29.93[†]	30.85[†]	19.43	6.63	3.64	12.47[†]
12	Sys. 8 → 10	-	-	-	-	38.96[†]	16.50[†]	29.77[†]	15.54[†]	27.58[†]	22.87[†]	20.66[†]	21.76[†]	10.41	3.88	3.02	7.04[†]

augmentation and speaker adaptation (Sys. 7 vs. 6). The AASR Conformer outperforms the comparable ASR baseline significantly by 1.19% absolute (8.67% relative) (Sys.10 vs. 9). Further with score fusion and cross system multi-pass decoding, the lowest overall WER of 7.04% is obtained (Sys. 12).

4.3. T-SNE Visualization And RMSE

Further analysis over the phonetic discrimination brought by UTI articulatory features via t-distributed stochastic neighbour embedding (t-SNE) [44] visualization is shown in the Figure 2 across four datasets. As expected, the original UTI articulatory features produce clear phonetic discrimination between example phonemes /t/ and /d/ on the TaL test set. Similar strong phonetic discrimination obtained using the cross-domain and cross-lingual generated UTI features can also be found in Figure 2 (b), (c) and (d) for the same pair of phonemes on the English DementiaBank Pitt, the Cantonese JCCOCC MoCA⁵ and the English TORGO test data. In addition, comparable RMSE scores (1.027, 1.030, 1.025 and 1.028) were obtained when performing in-domain and cross-domain UTI inversion for the TaL, DementiaBank Pitt, JCCOCC MoCA and TORGO corpora respectively. This demonstrates the robustness of our cross-domain and cross-lingual UTI inversion approach. In Figure 3, a further t-SNE plot of the generated articulatory features computed over phonemes /b/ and /d/ on the TORGO data with and without MLAN illustrates the importance of cross-domain and cross-lingual adaptation in UTI inversion.

5. Conclusion

This paper presents a cross-domain and cross-lingual A2A inversion approach that utilizes the parallel audio and ultrasound tongue imaging (UTI) data of the TaL corpus in A2A model training before being cross-domain and cross-lingual adapted to three datasets across two languages: English DementiaBank Pitt and Cantonese JCCOCC MoCA elderly speech corpora; and English TORGO dysarthric speech data to produce UTI based articulatory features. On three tasks, incorporating the A2A generated articulatory features consistently and significantly outperforms the baseline hybrid TDNN and end-to-end Conformer ASR systems constructed using acoustic features only. The proposed cross-domain and cross-lingual A2A inversion method allows more practical and wider use of UTI based

⁵UTI features are computed from tonal Cantonese phonemes with identical base phoneme but different tones are averaged.

articulatory features in elderly and disordered ASR systems.

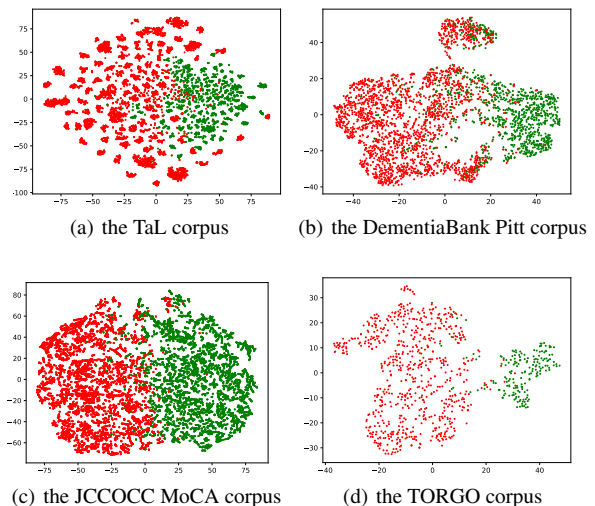


Figure 2: T-SNE plot of original (a) and cross-domain/lingual generated ((b) to (d)) UTI articulatory features computed over phonemes /t/ (in red) and /d/ (in green) on the TaL, Dementia-Bank Pitt, JCCOCC MoCA and TORGO corpora.




Figure 3: T-SNE plot of generated UTI articulatory features computed over phonemes /b/ (in red) and /d/ (in green) on TORGO data w/o cross-domain and cross-lingual adaptation.

6. Acknowledgement

This research is supported by Hong Kong RGC GRF grant No. 14200220, 14200021, TRS T45-407/19N, Innovation Technology Fund grant No. ITS/218/21, National Natural Science Foundation of China (NSFC) Grant 62106255 and Youth Innovation Promotion Association CAS Grant 2023119.

7. References

- [1] J. Yu, X. Xie *et al.*, “Development of the CUHK Dysarthric Speech Recognition System for the UASpeech Corpus.” in *INTERSPEECH*, 2018, pp. 2938–2942.
- [2] F. Xiong, J. Barker *et al.*, “Source domain data selection for improved transfer learning targeting dysarthric speech recognition,” in *ICASSP*. IEEE, 2020, pp. 7424–7428.
- [3] S. Liu *et al.*, “Recent Progress in the CUHK Dysarthric Speech Recognition System,” *TASLP*, vol. 29, pp. 2267–2281, 2021.
- [4] Z. Ye, S. Hu *et al.*, “Development of the CUHK Elderly speech recognition system for neurocognitive disorder detection using the dementiaabank corpus,” in *ICASSP*. IEEE, 2021, pp. 6433–6437.
- [5] M. Geng, X. Xie *et al.*, “Speaker adaptation using spectro-temporal deep features for dysarthric and elderly speech recognition,” *TASLP*, vol. 30, pp. 2597–2611, 2022.
- [6] Z. Yue, E. Loweimi *et al.*, “Acoustic Modelling From Raw Source and Filter Components for Dysarthric Speech Recognition,” *TASLP*, vol. 30, pp. 2968–2980, 2022.
- [7] S. Hu, X. Xie *et al.*, “Exploring Self-supervised Pre-trained ASR Models For Dysarthric and Elderly Speech Recognition,” in *ICASSP*. IEEE, 2023, pp. 1–5.
- [8] C. P. Ferri *et al.*, “Global prevalence of dementia: a Delphi consensus study,” *The lancet*, vol. 366, pp. 2112–2117, 2005.
- [9] F. Rudzicz *et al.*, “Speech recognition in Alzheimer’s disease with personal assistive robots,” in *SLPAT*, 2014, pp. 20–28.
- [10] L. Zhou *et al.*, “Speech Recognition in Alzheimer’s Disease and in its Assessment,” in *INTERSPEECH*, 2016, pp. 1948–1952.
- [11] B. Mirheidari, D. Blackburn *et al.*, “Dementia detection using automatic analysis of conversations,” *Computer Speech & Language*, vol. 53, pp. 65–79, 2019.
- [12] I. Kodrasi and H. Bourlard, “Spectro-temporal sparsity characterization for dysarthric speech detection,” *TASLP*, vol. 28, pp. 1210–1222, 2020.
- [13] B. L. Smith *et al.*, “Temporal characteristics of the speech of normal elderly adults,” *JSLHR*, vol. 30, pp. 522–529, 1987.
- [14] I. Zlokarnik, “Adding articulatory features to acoustic features for automatic speech recognition,” *JASA*, pp. 3246–3246, 1995.
- [15] K. Kirchhoff, G. A. Fink *et al.*, “Combining acoustic and articulatory feature information for robust speech recognition,” *Speech Communication*, vol. 37, pp. 303–319, 2002.
- [16] V. Mitra, G. Sivaraman *et al.*, “Hybrid convolutional neural networks for articulatory and acoustic information based speech recognition,” *Speech Communication*, vol. 89, pp. 103–112, 2017.
- [17] F. Rudzicz, “Articulatory knowledge in the recognition of dysarthric speech,” *TASLP*, vol. 19, pp. 947–960, 2010.
- [18] J. A. Gonzalez, L. A. Cheah *et al.*, “Direct speech reconstruction from articulatory sensor data by machine learning,” *TASLP*, vol. 25, pp. 2362–2374, 2017.
- [19] F. Xiong, J. Barker *et al.*, “Deep learning of articulatory-based representations and applications for improving dysarthric speech recognition,” in *ITG-Symposium*. VDE, 2018, pp. 1–5.
- [20] E. Yilmaz, V. Mitra *et al.*, “Articulatory and bottleneck features for speaker-independent ASR of dysarthric speech,” *Computer Speech & Language*, vol. 58, pp. 319–334, 2019.
- [21] O. Engwall, “Are static MRI measurements representative of dynamic speech? Results from a comparative study using MRI, EPG and EMA,” in *ICSLP*, 2000.
- [22] S. Narayanan, A. Toutios *et al.*, “Real-time magnetic resonance imaging and electromagnetic articulography database for speech production research (TC),” *JASA*, vol. 136, pp. 1307–1311, 2014.
- [23] M. Stone, “A guide to analysing tongue motion from ultrasound images,” *Clinical linguistics & phonetics*, vol. 19, pp. 455–501, 2005.
- [24] M. S. Ribeiro, A. Eshky *et al.*, “Speaker-independent classification of phonetic segments from raw ultrasound in child speech,” in *ICASSP*. IEEE, 2019, pp. 1328–1332.
- [25] J. Cleland, J. M. Scobbie *et al.*, “Enabling new articulatory gestures in children with persistent speech sound disorders using ultrasound visual biofeedback,” *JSLHR*, vol. 62, pp. 229–246, 2019.
- [26] M. S. Ribeiro, J. Sanger *et al.*, “TaL: a synchronised multi-speaker corpus of ultrasound tongue imaging, audio, and lip videos,” in *SLT*. IEEE, 2021, pp. 1109–1116.
- [27] J. Cai *et al.*, “Recognition and Real Time Performances of a Lightweight Ultrasound Based Silent Speech Interface Employing a Language Model,” in *INTERSPEECH*, 2011, pp. 1005–1008.
- [28] A. Eshky, M. S. Ribeiro *et al.*, “UltraSuite: a repository of ultrasound and acoustic data from child speech therapy sessions,” in *INTERSPEECH*, 2018, pp. 1888–1892.
- [29] G. Papcun, J. Hochberg *et al.*, “Inferring articulation and recognizing gestures from acoustics with a neural network trained on x-ray microbeam data,” *JASA*, vol. 92, pp. 688–700, 1992.
- [30] B. Uria, I. Murray *et al.*, “Deep architectures for articulatory inversion,” in *INTERSPEECH*, 2012, pp. 867–870.
- [31] X. Xie, X. Liu *et al.*, “Investigation of stacked deep neural networks and mixture density networks for acoustic-to-articulatory inversion,” in *ICSLP*. IEEE, 2018, pp. 36–40.
- [32] S. K. Maharana, A. Illa *et al.*, “Acoustic-to-Articulatory Inversion for Dysarthric Speech by Using Cross-Corpus Acoustic-Articulatory Data,” in *ICASSP*. IEEE, 2021, pp. 6458–6462.
- [33] D. Porras *et al.*, “DNN-based acoustic-to-articulatory inversion using ultrasound tongue imaging,” in *IJCNN*, 2019, pp. 1–8.
- [34] J. T. Becker, F. Boiler *et al.*, “The natural history of Alzheimer’s disease: description of study cohort and accuracy of diagnosis,” *Archives of neurology*, vol. 51, pp. 585–594, 1994.
- [35] S. S. Xu, M.-W. Mak *et al.*, “Speaker Turn Aware Similarity Scoring for Diarization of Speech-Based Cognitive Assessments,” in *APSIPA ASC*. IEEE, 2021, pp. 1299–1304.
- [36] F. Rudzicz, A. K. Namasivayam *et al.*, “The TORGO database of acoustic and articulatory speech from speakers with dysarthria,” *Language Resources and Evaluation*, vol. 46, pp. 523–541, 2012.
- [37] T. Hueber, E.-L. Benaroya *et al.*, “Development of a silent speech interface driven by ultrasound and optical images of the tongue and lips,” *Speech Communication*, vol. 52, pp. 288–300, 2010.
- [38] Y. Ji, L. Liu *et al.*, “Updating the silent speech challenge benchmark with deep learning,” *Speech Communication*, vol. 98, pp. 42–50, 2018.
- [39] M. S. Ribeiro, A. Eshky *et al.*, “Ultrasound Tongue Imaging for Diarization and Alignment of Child Speech Therapy Sessions,” in *INTERSPEECH*, 2019, pp. 16–20.
- [40] V. Peddinti, D. Povey *et al.*, “A time delay neural network architecture for efficient modeling of long temporal contexts,” in *INTERSPEECH*, 2015, pp. 3214–3218.
- [41] A. Gulati, J. Qin *et al.*, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *INTERSPEECH*, 2020, pp. 5036–5040.
- [42] P. Swietojanski, J. Li *et al.*, “Learning hidden unit contributions for unsupervised acoustic model adaptation,” *TASLP*, vol. 24, pp. 1450–1463, 2016.
- [43] M. Cui *et al.*, “Two-pass Decoding and Cross-adaptation Based System Combination of End-to-end Conformer and Hybrid TDNN ASR Systems,” in *INTERSPEECH*, 2022, pp. 3158–3162.
- [44] L. Van der Maaten and G. Hinton, “Visualizing data using t-SNE,” *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [45] P. Liu, Q. Yu *et al.*, “A deep recurrent approach for acoustic-to-articulatory inversion,” in *ICASSP*. IEEE, 2015, pp. 4450–4454.
- [46] S. Hu, S. Liu *et al.*, “Exploiting Cross Domain Acoustic-to-articulatory Inverted Features for Disordered Speech Recognition,” in *ICASSP*. IEEE, 2022, pp. 6747–6751.
- [47] P. Bell *et al.*, “Transcription of multi-genre media archives using out-of-domain data,” in *SLT*. IEEE, 2012, pp. 324–329.
- [48] M. Geng, X. Xie *et al.*, “Investigation of Data Augmentation Techniques for Disordered Speech Recognition,” in *INTERSPEECH*, 2020, pp. 696–700.
- [49] D. Povey, A. Ghoshal *et al.*, “The Kaldi speech recognition toolkit,” in *ASRU*. IEEE, 2011.
- [50] S. Watanabe, T. Hori *et al.*, “Espnet: End-to-end speech processing toolkit,” *arXiv:1804.00015*, 2018.