

THE SECRET SOURCE : INCORPORATING SOURCE FEATURES TO IMPROVE ACOUSTIC-TO-ARTICULATORY SPEECH INVERSION

Yashish M. Siriwardena and Carol Espy-Wilson

Institute for Systems Research, University of Maryland College Park, USA

ABSTRACT

In this work, we incorporated acoustically derived source features, aperiodicity, periodicity and pitch as additional targets to an acoustic-to-articulatory speech inversion (SI) system. We also propose a Temporal Convolution based SI system, which uses auditory spectrograms as the input speech representation, to learn long-range dependencies and complex interactions between the source and vocal tract, to improve the SI task. The experiments are conducted with both the Wisconsin X-ray microbeam (XRMB) and Haskins Production Rate Comparison (HPRC) datasets, with comparisons done with respect to three baseline SI model architectures. The proposed SI system with the HPRC dataset gains an improvement of close to 28% when the source features are used as additional targets. The same SI system outperforms the current best performing SI models by around 9% on the XRMB dataset.

Index Terms— source features, speech inversion, vocal tract variables, TCN, auditory spectrograms

1. INTRODUCTION

Speech production is a complex process involving coordinated movements of the articulators: lips, jaw, tongue, teeth, glottis and soft palate. The vibration of vocal folds at the glottis or the lack of it determines the periodicity of the generated sound. Apart from the vocal fold vibration, aspiration, frication and transients are identified as other sound sources [1]. The vocal tract which consists of the velum, tongue, lips and teeth, acts like an acoustic tube which modulates the source waveform. The acoustics of all vowel productions and most of the consonants have been described by a linear source-filter theory [1]. This theory is based on the assumption that the source of speech production is independent of the vocal tract filter. However, the actual process of speech production is nonlinear since the aerodynamics inside the glottis and vocal tract is governed by non-linear equations, and most importantly it has been shown that there exists a mutual interaction between the source and filter in certain cases of speech production [2]. In this work, we take into account this source-filter interaction to improve the acoustic-to-articulatory speech inversion task.

The inverse problem of determining the trajectories of the movement of speech articulators from the speech signal is referred to as acoustic-to-articulatory speech inversion [3, 4]. Learning this mapping from acoustics to articulation is an ill-posed problem which is known to be highly non-linear and non-unique [5]. Speaker variability in speech production makes it even harder or if not impossible

to develop speaker-independent SI systems. However, accurate estimation of articulatory trajectories or vocal tract variables can benefit speech applications like Automatic Speech Recognition (ASR) [6], speech synthesis [7], speech therapy [8] and mental health assessments [9, 10]. Ground-truth articulatory data are collected by techniques like X-ray microbeam [11], Electromagnetic Articulometry (EMA) [12] and real-time Magnetic Resonance Imaging (rt-MRI) [13]. All these methods are expensive, time consuming and need specialized equipment for observing articulatory movements directly [3]. For these reasons, developing a speaker-independent SI system that can accurately estimate articulatory features for any unseen speaker can potentially transform how speech research is conducted.

Recent advancements in deep neural networks (DNNs) and learning algorithms have significantly contributed to improving SI systems over the last few years. Bidirectional LSTMs (BiLSTMs) [14], BiGRNNs [15], CNN-BiLSTMs [16], Temporal Convolutional Networks (TCN) [17] and transformer models [18] have gained state-of-the-art results with multiple articulatory datasets. Most of these SI systems use either extracted acoustic features like Mel Frequency Cepstral Coefficients (MFCCs), Mel-spectrograms or the waveform itself as the input speech representation, and learns a mapping to the ground-truth articulatory variables. To further improve the SI task, phoneme features have been used as inputs along with the acoustic features in [16]. One limitation of these models is that you need phonetic transcriptions of the speech utterances at the time of inference. To address this issue while also leveraging additional information that phonetic transcriptions offer, multi task learning frameworks were proposed in [19] where phoneme labels are jointly predicted with TVs as targets. But, to the best of our knowledge, source level features have not been explored either as inputs or as targets to improve the SI task. One main reason for that can be the lack of ground-truth source data (like Electroglottography (EGG) recorded synchronously with EMA or XRMB recordings) that can be directly used to capture the glottal activity in the existing articulatory datasets. Hence, the current state of the art SI systems trained with these existing articulatory datasets are usually not able to estimate any of the glottal (or velar) activity in speech.

In this work, we investigated the idea of estimating source level features along with the TVs as targets to leverage any source-filter interactions to improve the overall SI task. Here we used source features from the Aperiodicity, Periodicity and Pitch (APP) detector [20] as proxies for the source activity. We also experimented with multiple input representations of speech with different DNN model architectures and used two publicly available articulatory datasets to show the significance of learning source level information in improving the acoustic-to-articulatory speech inversion task.

© 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

This work was supported by the National Science Foundation grant IIS1764010.

2. ARTICULATORY DATASETS

2.1. X-Ray Microbeam (XRMB) dataset

The original University of Wisconsin XRMB database [21] comprises of naturally spoken isolated sentences and short read paragraphs collected from 32 male and 25 female subjects. These speech utterances were recorded along with trajectory data captured by X-ray microbeam cinematography of the midsagittal plane of the vocal tract using pellets placed on several articulators: upper (UL) and lower (LL) lip, tongue tip (T1), tongue blade (T2), tongue dorsum (T3), tongue root (T4), mandible incisor (MANi), and (parasagittally placed) mandible molar (MANm). However, some of the articulatory recordings were marked as mistracked in the database and eliminating these samples left us with 46 speakers (21 males and 25 females) with a total of around 4 hours of speech data.

For each of the above mentioned articulators, the X-Y positions of the pellet movement is recorded. Since the X-Y positions of the pellets strongly depend on the anatomy of the speakers and variability of pellet placements, the measurements can vary significantly across speakers. Hence, to better represent vocal tract shape, relative measures were used to calculate the Tract Variables (TVs) from the X-Y positions of the pellets. TVs lead to a relatively speaker independent representation of speech articulation and characterize salient features of the vocal tract area function [22]. The TVs are based on articulatory phonology, a theoretical framework for speech production [23]. Using geometric transformations, the XRMB trajectories were converted to TV trajectories as outlined in [24]. The transformed XRMB database comprises of six TV trajectories: Lip Aperture (LA), Lip Protrusion (LP), Tongue Body Constriction Location (TBCL), Tongue Body Constriction Degree (TBCD), Tongue Tip Constriction Location (TTCL) and, Tongue Tip Constriction Degree (TTCD).

2.2. Haskins Production Rate Comparison (HPRC) dataset

We also used the HPRC database which contains recordings from 4 female and 4 male subjects reciting 720 phonetically balanced IEEE sentences [25] at normal and fast production rates [12]. The recordings were done using a 5-D electromagnetic articulometry (EMA) system (WAVE; Northern Digital). First, every sentence was produced at the speaker’s preferred ‘normal’ speaking rate and then a ‘fast’ repetition of the same, without making errors. Sensors were placed on the tongue (tip (TT), body (TB), root (TR)), lips (upper (UL) and lower (LL)) and mandible, together with reference sensors on the left and right mastoids, and upper and lower incisors (UI, LI). These EMA trajectories were obtained at 100 Hz and the synchronized audio were recorded at 44.1 KHz. Geometric transformations as defined in [3] were used to obtain 9 TVs (namely Lip Aperture (LA), Lip Protrusion (LP), Tongue Body Constriction Location (TBCL), Tongue Body Constriction Degree (TBCD), Tongue Tip Constriction Location (TTCL), Tongue Tip Constriction Degree (TTCD), Jaw Angle (JA), Tongue Middle Constriction Location (TMCL) and Tongue Middle Constriction Degree (TMCD)).

3. PROPOSED SPEECH INVERSION SYSTEM

3.1. Input representations

All the audio files are first segmented into 2second long segments and the shorter audios are zero padded at the end. However for the HPRC dataset, the audio files are first down-sampled to 16 KHz before segmentation and padding. The following features are then extracted from the segmented audio utterances.

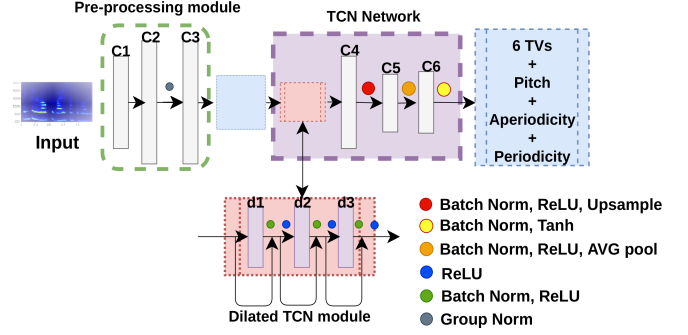


Fig. 1: Model architecture of the SI system. Here C1-C6 represent 1D-CNN layers where as d1-d3 represent 1D dilated CNN layers

3.1.1. Auditory Spectrograms

We converted the one-dimensional pressure time waveform into a two-dimensional pattern of neural activity distributed along the tonotopic axis (roughly a logarithmic frequency). This two-dimensional representation, which is defined as an ‘auditory spectrogram’ (Audspec) [26] is used in our proposed SI system as the input speech representation. It has been shown that this Audspec is an enhanced and a noise-robust estimate of the Fourier-based spectrogram [26].

3.1.2. MFCCs and Melspectrograms

Mel-Frequency Cepstral Coefficients (MFCCs) and Melspectrograms (MSPECs) are extracted as the acoustic input features for baseline SI systems. Both MFCCs and MSPECs were extracted using a 20ms Hamming analysis window with a 10ms frame shift. For MFCCs, 13 cepstral coefficients were extracted for each frame while 40 Mel frequencies were used for both MFCCs and MSPECs. Both MFCCs and MSPECs are utterance wise normalized (z-normalized) prior to model training.

3.2. Model Architecture and training

We propose a Temporal Convolution Network (TCN) based acoustic-to-articulatory SI system which takes in the Audspecs as input. The proposed system, unlike conventional SI systems, estimates both TVs and source level features (aperiodicity, periodicity and pitch) as the output. The model is optimized using the Mean Squared Error (MSE) loss computed between the predicted articulatory variables and the ground truth (TVs from articulatory datasets and source features from APP detector [20]).

The SI system is implemented in PyTorch with 1-D convolutional (CNN) layers. The complete network is inspired by the multi-layered Temporal Convolution Network in [27]. Figure 1 shows the proposed model architecture with its sub-modules used for pre processing and dilated TCN. The Pre-processing module contains three 1-D CNN layers with 1×1 kernels (C1, C2 and C3), which have 128, 256 and 256 filters, respectively. The d1, d2 and d3 dilated CNN layers have a kernel size of 3 with 1,4 and 16 dilation rates respectively. Upsampling (window size 4) is done after C4 layer and average pooling (window size 5) is done after C5 layer along with BatchNorm layers after every CNN layer in the TCN network. The upsampling and average pooling operations take care of matching the time dimension of the input spectrograms to the target time dimension of TVs.

To train the SI system, learning rates were determined based on a grid search by testing all combinations from [1e-2, 1e-3, 1e-4, 3e-4] that resulted in 1e-3 as the best pick. A similar grid search

Table 1: PPMC scores for articulatory variable prediction on the XRMB dataset. Model names with ‘SF’ uses source features as additional targets. The AVG. TVs column for those models also show the percentage increase in TV prediction with respect to the same model which does not use source features

Model	LA	LP	TBCL	TBCD	TTCL	TTCD	Ap.	Per.	Pitch	AVG. TVs	AVG. all
TCN-Audspec	0.7977	0.7942	0.7883	0.7836	0.7743	0.7684	-	-	-	0.7844	-
TCN-SF-Audspec	0.8448	0.8640	0.8604	0.8818	0.9029	0.9005	0.9082	0.8860	0.9021	0.8770 (9.3%)	0.8834
TCN-Mspec	0.7432	0.7427	0.7366	0.7244	0.7179	0.6993	-	-	-	0.7273	-
TCN-SF-Mspec	0.8364	0.8639	0.8727	0.8607	0.8807	0.8917	0.8732	0.9005	0.8638	0.8677 (14%)	0.8715
BiGRNN-MFCC	0.8801	0.6200	0.8580	0.7382	0.6922	0.9206	-	-	-	0.7848	-
BiGRNN-SF-MFCC	0.8810	0.6211	0.8628	0.7365	0.7019	0.9191	0.8693	0.9163	0.7209	0.7871 (0.2%)	0.8032
CNN-BiGRNN-Mspec	0.8801	0.6165	0.8505	0.7355	0.7146	0.9171	-	-	-	0.7858	-
CNN-BiGRNN-SF-Mspec	0.8799	0.6246	0.8566	0.7302	0.7065	0.9175	0.8794	0.9296	0.7441	0.7859 (0.01%)	0.8076
CNN-BLSTM-Mspec	0.8770	0.6184	0.8463	0.7200	0.6915	0.9197	-	-	-	0.7788	-
CNN-BLSTM-SF-Mspec	0.8774	0.6202	0.8525	0.7172	0.6941	0.9180	0.8734	0.9263	0.7442	0.7799 (0.1%)	0.8026

was done to choose the batch size from [16, 32, 64, 128] and 64 gave the best validation MSE. The objective function was optimized using the ADAM optimizer with an ‘ExponentialLR’ learning rate scheduler and a decay of 0.5. All models were trained by monitoring the validation loss.

To train the models with the XRMB dataset, the dataset was divided into training, development, and testing splits, so that the training set has utterances from 36 speakers and the development and testing sets have 5 speakers each (3 males, 2 females). To train the models with the HPRC dataset, similar to the XRMB dataset, the dataset was divided into training, development, and testing sets, so that the training set has utterances from 6 speakers (3 Males, 3 Females) and the development and testing sets have utterances of 2 speakers (1 male, 1 female) equally split between them. We used audio samples with both the normal and fast production rates in the HPRC dataset. To create the validation and test sets, the audio samples from the two speakers were randomly split, so that both the splits have samples corresponding to normal and fast rates.

For both the datasets, none of the training splits have overlapping speakers with the development and testing sets and hence all the models are trained in a ‘speaker-independent’ fashion. The splits also ensured that around 80% of the total number of utterances were present in training, and the development and testing sets have a nearly equal number of utterances. This allocation was done in a completely random manner.

3.3. Performance metric

All the models are evaluated with the Pearson Product Moment Correlation (PPMC) scores computed between the estimated articulatory variables and the corresponding ground-truth. Equation 1 is used to compute the PPMC score, where X represents the estimated articulatory variable, \bar{X} the mean of the estimated, Y the ground-truth variable, \bar{Y} the mean of the ground-truth variable and N the number of articulatory variables.

$$PPMC = \frac{\sum_i^N (X[i] - \bar{X})(Y[i] - \bar{Y})}{\sqrt{\sum_i^N (X[i] - \bar{X})^2 (Y[i] - \bar{Y})^2}} \quad (1)$$

4. BASELINE SPEECH INVERSION SYSTEMS

This section discusses about the baseline SI systems implemented for comparison. Detailed information on the model architectures and implementation can be found in a GitHub repository ¹

¹<https://github.com/Yashish92/Speech-Inversion-TCN>

Table 2: PPMC scores for HPRC dataset.

Model	AVG. 9 TVs	AVG. all
TCN-Audspec	0.4805	-
TCN-SF-Audspec	0.7573 (27.7%)	0.7636
TCN-Mspec	0.4763	-
TCN-SF-Mspec	0.6503 (17.4%)	0.6621
BiGRNN-MFCC	0.7118	-
BiGRNN-SF-MFCC	0.7153 (0.3%)	0.7263
CNN-BiGRNN-Mspec	0.7277	-
CNN-BiGRNN-SF-Mspec	0.7290 (0.1%)	0.7461
CNN-BLSTM-Mspec	0.7245	-
CNN-BLSTM-SF-Mspec	0.7259 (0.1%)	0.7428

4.1. BiGRNN model

We used the BiGRNN model architecture implemented in [15, 19] as one of our baseline models for comparison. The model has 2 bidirectional layers of Gated Recurrent Units (GRUs) followed by two time distributed fully connected layers. Dropout layers are also used after every layer to minimize the issue of over-fitting. 13 MFCCs (extracted as in section 3.1.2) are used as input to this SI system.

4.2. CNN-BiLSTM and CNN-BiGRNN model

A CNN-BiLSTM model inspired by the work in [16] was implemented as another baseline SI system. The model consists of 5 CNN (1D CNNs) layers, whose outputs are then concatenated together and fed to 2 BiLSTM layers. The output from the last BiLSTM layer is then passed through two time distributed fully connected layers, where the final fully connected layer serves as the output layer.

A similar architecture was used to implement a CNN-BiGRNN model where the only difference is that the 2 BiLSTM layers in the CNN-BiLSTM model are now replaced with bidirectional layers of Gated Recurrent Units (GRUs). CNN-BiGRNN model is comparatively light weight due to the GRUs used in the model instead of LSTM units. A BiGRNN based SI system has also shown to outperform a conventional BiLSTM based SI system in [15] which motivated this new CNN-BiGRNN model as a baseline for comparison.

Similar to the BiGRNN model in section 4.1, dropout layers are used after every layer to minimize possible over-fitting in both the CNN-BiLSTM and CNN-BiGRNN models. Both models used MSPECs as the input speech representation.

5. RESULTS

5.1. Comparison with baseline SI systems

The baseline models discussed in section 4 were trained and evaluated with the same train-dev-test splits for comparison. For every

model architecture, two versions of the model were implemented with one only predicting the TVs as targets (6 TVs for XRMB dataset, 9 TVs for HPRC dataset) and the other predicting both TVs and source features.

A baseline TCN based SI system (TCN-Audspec) was trained with Audspecs as input and only TVs as targets for both XRMB and HPRC datasets. A similar TCN architecture was implemented to use MSPECs as inputs and, two versions of this model (TCN-Mspec and TCN-SF-Mspec) were trained similar to the other baseline models. Table 1 shows the PPMC scores for TV estimation on the XRMB dataset. The PPMC scores for individual TVs and source features are listed here along with average scores across TVs and all the predicted articulatory variables (TVs + source features). Similarly, Table 2 lists the average PPMC scores across the 9 TVs and all the articulatory variables for the HPRC dataset.

5.2. Estimated TVs and source features

Figure 2 shows the estimated constriction degree TVs (LA, TBCD, TTCD) and source features from the proposed TCN-SF-Audspec and the TCN-Audspec models. As can be observed in the plots, the source features are predicted with a considerably better accuracy, which is an added advantage of the proposed SI system. This also gives an almost complete articulatory representation of speech (only missing velar activity) that can be useful in various speech applications (e.g articulatory speech synthesis).

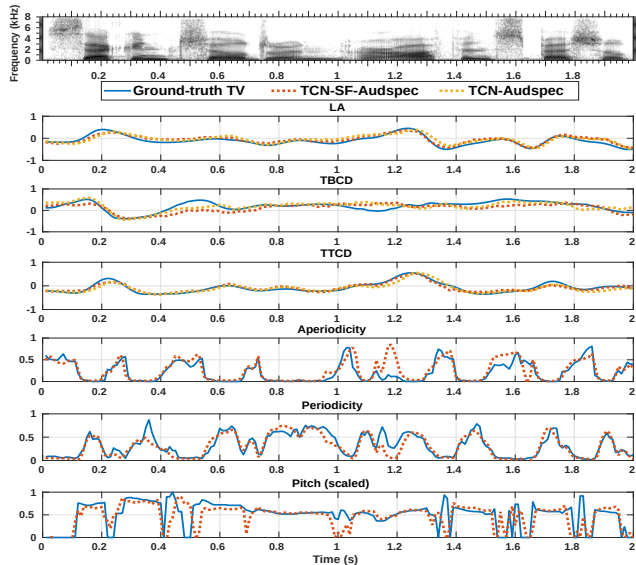


Fig. 2: LA and constriction degree TVs + source features for the utterance ‘second children are often special’ estimated by the proposed TCN-SF-Audspec model compared to the TCN-Audspec. Solid blue Line - ground truth, red dotted line - predictions by the TCN-SF-Audspec, yellow dotted Line - predictions by TCN-Audspec.

6. DISCUSSION

Deep Neural Networks have been effective in learning complex non-linear relationships between input speech representations and articulatory movements that has led to the recent success in acoustic-to-articulatory SI task. However, there is still a lot of room for improvement, especially in developing speaker-independent and more generalizable SI systems that can be effectively utilized in speech applications. To this end, we explored the idea of incorporating source

characteristics to exploit any source-filter interactions, and thereby mutually learn, and improve, the acoustic-to-articulatory SI task.

We used aperiodicity, periodicity and pitch as source features, which are used as additional targets to train the SI systems. For the proposed TCN based SI system and for every baseline model, two versions of the model were trained with the goal of investigating the real effect of incorporating source features. As shown in Tables 1 and 2, the results are consistent across both articulatory datasets, and support the fact that incorporating source features into the mix of TVs is definitely helping the estimation of articulatory variables. This observation is quite evident with the proposed TCN models which use Audspecs or MSPECs as inputs. For example, with the XRMB dataset, the TCN model which uses MSPECs gain an absolute improvement of around 14% with respect to the same model which does not use source features as targets. Similarly with the HPRC dataset, the TCN model which uses Audspecs gain an absolute improvement of close to 28% when the source features are used as additional targets. Most importantly, when the best PPMC scores for average TV estimation with the proposed TCN model is considered, it is around a 9% improvement over the current best performing SI systems in [3, 15], trained and evaluated on the same splits of the XRMB dataset.

However, a key observation here is that both the input speech representation and the DNN model architecture play a significant role in learning these complex relationship between the source features and TVs. For example, with the 13 MFCCs, which is the most commonly used speech representation in SI systems, adding source features as targets does not significantly improve the PPMC scores. This is consistent with both the XRMB and HPRC datasets with the BiGRNN model which uses 13 MFCCs as input. This can be mainly due to the fact that the 13 MFCCs do not contain important source information and is usually limited to capturing the filter characteristics (vocal tract) in speech production. Moreover, having richer speech representations that contain source information does not necessarily mean it will improve on the SI task. A fine observation to support that is the CNN-BiGRNN and CNN-BiLSTM models that use MSPECs as inputs, which necessarily contain valuable source information unlike the 13 MFCCs. This elucidate the fact that the DNN model architecture too plays a critical role in learning these complex dependencies between the source and articulatory targets.

TCN based models have shown to be extremely effective in speech applications [28] with learning long-range temporal (and contextual) dependencies. These models have shown to outperform typical RNN and CNN based models, especially in applications where subtle and complex contextual information needs to be extracted from input representations [27]. This was one of the key motivations for the proposed SI system, which ultimately outperformed all the other baseline systems on both XRMB and HPRC datasets. However, as mentioned earlier, the input representation used to train these models also play a role which can be clearly observed with results in Tables 1 and 2. The TCN model trained with Audspecs as input is outperforming the same TCN based model architecture trained with MSPECs which suggests that the Audspecs might be capturing important spectral and temporal information that is helpful in mutually learning both the source features and TVs.

Conventional SI systems usually predict constriction degree related TVs significantly better (eg., LA, TTCD) with respect to the constriction location related TVs (eg., LP, TTCL). The same can be more or less observed (with the exception of TBCL and TBCD) for individual TV predictions in Table 1 for all the ‘baseline models’. Surprisingly, this is not as evident with the TCN based models which tend to predict both the location and degree TVs with close

to similar accuracies. Moreover, further analysis needs to be done to investigate the ways and instances by which the source features are actually interacting with the TVs, and also to understand what the TCN models are actually capturing as source-filter interactions that is ultimately helping the overall SI task.

7. REFERENCES

- [1] Kenneth N Stevens, *Acoustic phonetics*, vol. 30, MIT press, 2000.
- [2] Ingo Titze, Tobias Riede, and Peter Popolo, “Nonlinear source-filter coupling in phonation: vocal exercises,” *J Acoust Soc Am*, vol. 123, no. 4, pp. 1902–1915, Apr. 2008.
- [3] Ganesh Sivaraman, Vikramjit Mitra, Hosung Nam, Mark Tiede, and Carol Espy-Wilson, “Unsupervised speaker adaptation for speaker independent acoustic to articulatory speech inversion,” *The Journal of the Acoustical Society of America*, vol. 146, no. 1, pp. 316–329, 2019.
- [4] Abdolreza Sabzi Shahrehabaki, Giampiero Salvi, Torbjørn Svendsen, and Sabato Marco Siniscalchi, “Acoustic-to-articulatory mapping with joint optimization of deep speech enhancement and articulatory inversion models,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 135–147, 2022.
- [5] Chao Qin and Miguel Á Carreira-Perpiñán, “An empirical investigation of the nonuniqueness in the acoustic-to-articulatory mapping,” *Interspeech*, pp. 74–77, 2007.
- [6] Vikramjit Mitra, Hosung Nam, Carol Y. Espy-Wilson, Elliot Saltzman, and Louis Goldstein, “Retrieving tract variables from acoustics: A comparison of different machine learning strategies,” *IEEE Journal on Selected Topics in Signal Processing*, vol. 4, no. 6, pp. 1027–1045, sep 2010.
- [7] Zhen-Hua Ling, Korin Richmond, and Junichi Yamagishi, “Articulatory control of hmm-based parametric speech synthesis using feature-space-switched multiple regression,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 1, pp. 207–219, 2013.
- [8] Sascha Fagel and Katja Madany, “A 3-d virtual head as a tool for speech therapy for children,” in *INTERSPEECH*, 2008.
- [9] Carol Espy-Wilson, Adam C. Lammert, Nadee Seneviratne, and Thomas F. Quatieri, “Assessing Neuromotor Coordination in Depression Using Inverted Vocal Tract Variables,” in *Proc. Interspeech 2019*, 2019, pp. 1448–1452.
- [10] Yashish M. Siriwardena, Carol Espy-Wilson, Chris Kitchen, and Deanna L. Kelly, *Multimodal Approach for Assessing Neuromotor Coordination in Schizophrenia Using Convolutional Neural Networks*, p. 768–772, Association for Computing Machinery, New York, NY, USA, 2021.
- [11] John R Westbury, “Speech Production Database User’s Handbook,” *IEEE Personal Communications - IEEE Pers. Commun.*, vol. 0, no. June, 1994.
- [12] Mark Tiede, Carol Y. Espy-Wilson, Dolly Goldenberg, Vikramjit Mitra, Hosung Nam, and Ganesh Sivaraman, “Quantifying kinematic aspects of reduction in a contrasting rate production task,” *The Journal of the Acoustical Society of America*, vol. 141, no. 5, pp. 3580–3580, 2017.
- [13] Shrikanth Narayanan, Krishna Nayak, Sungbok Lee, Abhinav Sethy, and Dani Byrd, “An approach to real-time magnetic resonance imaging for speech production,” *The Journal of the Acoustical Society of America*, vol. 115, no. 4, pp. 1771–1776, mar 2004.
- [14] Aravind Illa and Prasanta Kumar Ghosh, “Low Resource Acoustic-to-articulatory Inversion Using Bi-directional Long Short Term Memory,” in *Proc. Interspeech 2018*, 2018, pp. 3122–3126.
- [15] Yashish M. Siriwardena, Ahmed Adel Attia, Ganesh Sivaraman, and Carol Espy-Wilson, “Audio data augmentation for acoustic-to-articulatory speech inversion using bidirectional gated rnns,” arXiv 2022.
- [16] Abdolreza Sabzi Shahrehabaki, Sabato Marco Siniscalchi, Giampiero Salvi, and Torbjørn Svendsen, “Sequence-to-Sequence Articulatory Inversion Through Time Convolution of Sub-Band Frequency Signals,” in *Proc. Interspeech 2020*, 2020, pp. 2882–2886.
- [17] Abdolreza Sabzi Shahrehabaki, Sabato Marco Siniscalchi, and Torbjørn Svendsen, “Raw Speech-to-Articulatory Inversion by Temporal Filtering and Decimation,” in *Proc. Interspeech 2021*, 2021, pp. 1184–1188.
- [18] Sathvik Udupa, Anwesa Roy, Abhayjeet Singh, Aravind Illa, and Prasanta Kumar Ghosh, “Estimating Articulatory Movements in Speech Production with Transformer Networks,” in *Proc. Interspeech 2021*, 2021, pp. 1154–1158.
- [19] Yashish M. Siriwardena, Ganesh Sivaraman, and Carol Espy-Wilson, “Acoustic-to-articulatory Speech Inversion with Multi-task Learning,” in *Proc. Interspeech 2022*, 2022, pp. 5020–5024.
- [20] O. Deshmukh, C. Y. Espy-Wilson, A. Salomon, and J. Singh, “Use of temporal information: detection of periodicity, aperiodicity, and pitch in speech,” *IEEE Transactions on Speech and Audio Processing*, vol. 13, no. 5, pp. 776–786, 2005.
- [21] John R Westbury, “Speech Production Database User’s Handbook,” *IEEE Personal Communications*, vol. 0, no. June, 1994.
- [22] Richard S. McGowan, “Recovering articulatory movement from formant frequency trajectories using task dynamics and a genetic algorithm: Preliminary model tests,” *Speech Communication*, vol. 14, no. 1, pp. 19–48, 1994.
- [23] Catherine P Browman and Louis Goldstein, “Articulatory Phonology : An Overview *,” *Phonetica*, vol. 49, pp. 155–180, 1992.
- [24] Vikramjit Mitra, Hosung Nam, Carol Espy-Wilson, Elliot Saltzman, and Louis Goldstein, “Recognizing articulatory gestures from speech for robust speech recognition,” *The Journal of the Acoustical Society of America*, vol. 131, no. 3, pp. 2270–2287, 2012.
- [25] “Ieee recommended practice for speech quality measurements,” *IEEE Transactions on Audio and Electroacoustics*, vol. 17, no. 3, pp. 225–246, 1969.
- [26] Kuansan Wang and S. Shamma, “Self-normalization and noise-robustness in early auditory representations,” *IEEE Transactions on Speech and Audio Processing*, vol. 2, no. 3, pp. 421–435, 1994.
- [27] Colin Lea, Michael D. Flynn, Rene Vidal, Austin Reiter, and Gregory D. Hager, “Temporal convolutional networks for

action segmentation and detection,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

- [28] Ashutosh Pandey and DeLiang Wang, “Tcnn: Temporal convolutional neural network for real-time speech enhancement in the time domain,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6875–6879.