



ACOUSTIC TO ARTICULATORY SPEECH INVERSION USING MULTI-RESOLUTION SPECTRO-TEMPORAL REPRESENTATIONS OF SPEECH SIGNALS

Rahil Parikh^{1*}, Nadee Seneviratne^{1*}, Ganesh Sivaraman², Shihab Shamma¹, Carol Espy-Wilson¹

¹Institute for Systems Research,
University of Maryland College Park, USA
²Pindrop Inc., USA

rahil@umd.edu, nadee@umd.edu

Abstract

Multi-resolution spectro-temporal features of a speech signal represent how the brain perceives sounds by tuning cortical cells to different spectral and temporal modulations. These features produce a higher dimensional representation of the speech signals. The purpose of this paper is to evaluate how well the auditory cortex representation of speech signals contribute to estimate articulatory features of those corresponding signals. Since obtaining articulatory features from acoustic features of speech signals has been a challenging topic of interest for different speech communities, we investigate the possibility of using this multi-resolution representation of speech signals as acoustic features. We used U. of Wisconsin X-ray Microbeam (XRMB) database of clean speech signals to train a feed-forward deep neural network (DNN) to estimate articulatory trajectories of six tract variables. The optimal set of multi-resolution spectro-temporal features to train the model were chosen using appropriate scale and rate vector parameters to obtain the best performing model. Experiments achieved a correlation of 0.675 with ground-truth tract variables. We compared the performance of this speech inversion system with prior experiments conducted using Mel Frequency Cepstral Coefficients (MFCCs).

Index Terms: spectro-temporal features, speech inversion, vocal tract variables, auditory cortex, deep neural networks

1. Introduction

Acoustic to articulatory speech inversion is the process of mapping the speech acoustic signal to articulatory parameters associated with the movement of the vocal tract articulators for speech production. When accurately estimated, articulatory movements of the vocal tract can be applied to speech accent conversion [1], speech therapy [2], language learning, Automatic Speech Recognition (ASR) [3], and detection of depression from speech.

Real articulatory data is obtained from subjects using techniques like Electromagnetic Articulometry (EMA), X-ray microbeam, and real-time Magnetic Resonance Imaging (rt-MRI). However these techniques are expensive and time consuming. Obtaining real articulatory data is not practically feasible for real world applications like ASR. Hence, it is essential to develop speech inversion systems that are speaker independent and can accurately estimate articulatory features from the speech signal for any unseen test speaker.

* These authors contributed equally to this work.

The mapping from acoustics to articulations is known to be highly non-linear and non-unique [4]. Adding speaker variability to the already challenging problem makes it even more difficult. Based on a comprehensive study of the speech inversion techniques, the speaker dependent techniques can be classified into three categories-

- Codebook based approaches [5, 6] in which a codebook of acoustic and corresponding articulatory patterns is constructed from the training data
- Analytical approaches involving articulatory models such as Maeda's [7]
- Statistical modeling (parametric and non-parametric) of acoustic to articulatory mapping like Gaussian Mixture Model (GMM), Mixture density networks (MDN) [8], Hidden Markov models (HMM) [9], generalized smoothness criteria [10], and neural networks [11]

There have been a few attempts to perform speaker independent speech inversion using deep neural networks (DNN) [12]. Previously we worked on improving the performance of this baseline DNN based speech inversion system in [12]. The first improvement was to make the system more robust to noise [13]. Secondly, we trained a multi-task learning based DNN to train the SI system on multiple articulatory databases in order to increase the generalizability of the TV estimation task [14].

Most of the approaches for speech inversion have parameterized the speech signal as mel frequency cepstral coefficients (MFCCs), mel filterbanks, perceptual linear prediction (PLP) coefficients which are commonly used acoustic features for ASR. The homomorphic filtering based cepstral coefficients make sense for speech inversion because they keep only the slowly changing content of the speech signal associated with articulation and discard the fast changing content associated with voice quality and identity. However, with the advent of deep and convolutional neural networks, richer and complete representation of the speech signal has been favored over features like cepstral coefficients.

The auditory cortex represents speech using multiple spectral and temporal resolutions using cells tuned to different spectral and temporal modulations operating through different frequency specific neural processing channels [15]. Thus, the brain performs a higher dimensional representation of the speech spectrum. In the neuroscience literature there are two contending theories of speech perception [16], namely - (1) General Auditory theory [17, 18] and (2) Motor theory of speech perception [19, 20]. The auditory theory argues that speech perception happens solely due to representations of speech in the auditory cortical regions. The motor theory says that speech production

information is essential in recognizing speech and contributes to the robustness of the perception. This study brings together the two theories by attempting to estimate articulatory movements using cortical-like multi-resolution spectro-temporal representations of speech. These multi-resolution spectro-temporal representations will be referred to as cortical features in this paper. The cortical features provide a fine grained control over the modulation rates to focus on in the representations. They also provide us valuable knowledge about the relative contributions of different spectral and temporal modulation scales and rates towards estimating the articulatory movements. This knowledge can potentially lead to more robust acoustic feature representations targeted towards optimally capturing the speech articulations while discarding non-speech auditory objects. To the best of our knowledge cortical features have not been used to perform acoustic-to-articulatory speech inversion.

In this paper, we focus on developing a speaker independent speech inversion system using multi-resolution spectro-temporal representations of natural articulatory speech data. We used the multi-speaker XRMB clean speech dataset to perform our experiments. Further details about this dataset is given in section 2. We used the Matlab toolbox developed by the U. of Maryland Neural Systems Laboratory to extract the cortical features of the speech signals. The extracted features' dimensionality was reduced using Higher Order Singular Value Decomposition (HOSVD) before passing to a simple feed-forward neural network to be trained. The development of the model is described in section 3. The following section 4 will explain the experiments conducted to obtain the optimum parameters and the best performing network along with the results obtained by comparing the model developed with similar other systems which were using MFCCs as acoustic features. Section 5 presents the conclusion.

2. Dataset Description

To train the model to estimate Tract Variables (TVs), it is necessary to have a dataset which has ground-truth TVs corresponding to its speech utterances. Thus, we used the Wisconsin X-Ray Microbeam database [21] to perform our experiments. The choice of this database also helps to compare the performance of this model with similar experiments conducted earlier using different types of acoustic features.

The XRMB recordings comprise of naturally spoken utterances from 32 male and 25 female subjects along with X-ray Microbeam cinematography of the mid-sagittal plane of the vocal tract with pellets placed at points along the vocal tract. The trajectory data are recorded for the individual articulators: Upper Lip, Lower Lip, Tongue Tip, Tongue Blade, Tongue Dorsum, Tongue Root, Lower Front Tooth (Mandible Incisor), Lower Back Tooth (Mandible Molar). We call these trajectories as pellet trajectories. The miss tracked segments were removed from the database before using it for our analysis.

The TVs specify the salient features of the vocal tract area function more directly than the pellet trajectories [22] and are relatively speaker independent. They also provide us a theoretical framework to analyze speech production with the theoretical framework of articulatory phonology. Hence, the pellet trajectories were converted to TV trajectories using geometric transformations as outlined in [23]. The transformed XRMB database consists of 21 males and 25 females, with a total of 4 hours of speech data with corresponding 6 TV trajectories. The TVs obtained from the seven pellet trajectories were – Lip Aperture (LA), Lip Protrusion (LP), Tongue Body Constriction Location

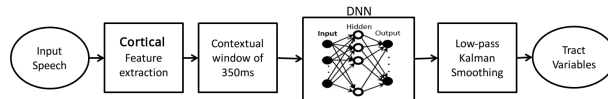


Figure 1: Block diagram of the speech inversion system

(TBCL), Tongue Body Constriction Degree (TBCD), Tongue Tip Constriction Location (TTCL) and, Tongue Tip Constriction Degree (TTCD).

3. Speech Inversion System

The steps involved in building the model depicted in Figure 1 to estimate TVs are explained in this section. The details of extracting multi-resolution spectro-temporal features of the speech signal, performing feature reduction on the same, and training and testing the DNN are elaborated below.

3.1. Feature Extraction

The cortical features or the multi-resolution spectro-temporal representations [24] are obtained using an existing toolbox [25] developed to perform these operations. The audio files are first converted to an auditory spectrogram to emulate the processing performed by the ear [24]. A bank of constant-Q bandpass filters are used to model the human cochlea. The NSL toolbox uses a cochlear filter bank of 128 filters with center frequencies uniformly distributed along the logarithmic axis, spread over 5.33 octaves. The hair cells and the lateral inhibitory network are modelled as explained in [26] to yield the auditory spectrogram. The cortical representations are obtained by performing a wavelet transform of this spectrogram with Spectro-Temporal Response Fields (STRFs), which are a bank of 2-Dimensional filters that are tuned to temporal modulations (or rates) which vary from slow to fast, and spectral modulations (or scales) that vary from narrow to broad [27]. This mimics the role of the primary auditory cortex in computing the temporal and spectral modulation content of the auditory spectrogram [24]. These are a three dimensional tensor for each frame of the input audio signal.

In our analysis, we generate the auditory spectrogram using a frame length of 10ms, to correspond to the sampling time of the tract variables signals. The cortical features are generated using STRFs with temporal modulation filters at 2,4,8,16 and 32 Hz for upward and downward moving patterns in the spectrogram, and spectral modulation filters at 1,2,4 and 8 cycles/octave. The features are thus a tensor with dimensions $4 \times 10 \times 128$ representing the 4 scale values, 10 rate values and 128 frequency channels, for each time frame of the input audio-sample.

3.2. Dimensionality Reduction

The cortical features span over 5120 dimensions for each time frame of 10ms, obtained by 4 scale filters, 10 rate filters, and 128 frequency channels. In such a high dimensional representation, the features are sparse and any machine learning algorithm will suffer from the curse of dimensionality. Fitting a model using such a high dimensional feature would require a huge number of parameters, making it computationally intensive. Hence an efficient dimension reduction method is required to reduce the dimensionality of the features. Unfortunately, conventional dimension reduction methods such as Principal Com-

ponent Analysis (PCA) fail to discriminate between the contributions of the various feature sub-spaces, such as scales, rates and frequency, and treat all modes similarly. Dimensionality reduction is thus obtained by using principal components from HOSVD [28]. Since we wish to retain all the information with respect to time, dimension reduction is only performed on the three dimensional scale-rate-frequency tensor. The HOSVD can be obtained by performing SVD on matrices obtained by unfolding the mode of the tensor along which the dimension needs to be reduced [29]. A tensor can be decomposed into another tensor of the same dimensions and the left-singular matrices obtained by unfolding each mode of the tensor, as shown in the equation below [28]:

$$T = A \times_1 U^s \times_2 U^r \times_3 U^f \quad (1)$$

where U^s , U^r , U^f , are the unitary left-singular matrices corresponding to the scales, rates and frequencies, obtained by unfolding each mode of the original tensor, and A is an ordered and orthogonal tensor. The tensor T in equation (1) is 3-Dimensional tensor, obtained by concatenating the absolute values of cortical features of all samples in the training set.

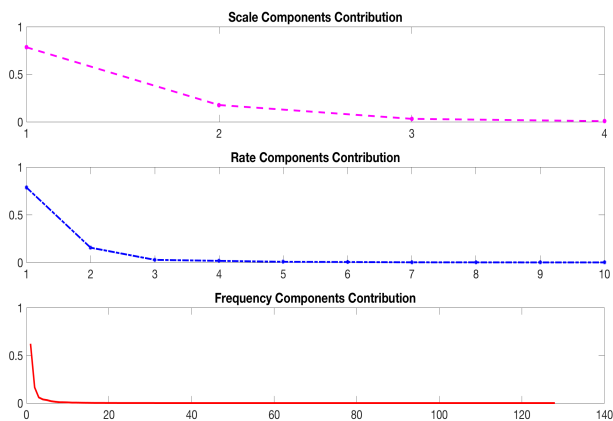


Figure 2: *Principal Components Contributions for Scales, Rates and Frequencies*

To determine the contribution of the principal components (PCs) for each mode, we compute the relative energy of each eigenvalue for the given mode, using the following formula:

$$\alpha_{i,j} = \frac{\lambda_{i,j}}{\sum_{k=1}^{N_i} \lambda_{i,k}} \quad (2)$$

where $\lambda_{i,j}$ is the corresponding eigenvalue of the j^{th} PC of the subspace S_i and N_i is the dimension of S_i . These plots are shown in Figure 2.

Components that contribute significantly are retained, and the tensor is multiplied to the corresponding eigenvectors for the given mode.

For instance, as shown in Figure 2 only the first 8 principal components contribute significantly to the information in the subspace. The matrix with eigenvectors corresponding to these 8 eigenvalues are multiplied with the tensor A of dimensions $4 \times 10 \times 128$, to obtain a tensor A_f with dimensions $4 \times 10 \times 8$. In our experiments, the input tensor of size $4 \times 10 \times 128$ is reduced to a tensor of size $4 \times 5 \times 7$ and size $4 \times 6 \times 8$ for each time frame of an audio-sample. This tensor is then vectorized for each time frame.

3.3. DNN Training

The input layer accepts contextualized cortical features and the output layer (which estimates the TVs) has a dimensionality of 6 nodes. The input dataset was divided into training, development, and testing sets so that the training set has utterances from 36 speakers and the development and testing sets have 5 speakers each (3 males, 2 females). Around 80% of the total number of utterances were present in the training set. The development and testing sets have a nearly equal number of utterances. This allocation of utterances was done in a completely random manner.

In this work we aim to demonstrate that cortical features can be mapped to tract variables even with minimal complexity of the neural network. We thus trained three feed-forward DNNs with 400, 512, and 800 neurons each in all the hidden layers. Similar to prior experiments [30], we observed that as the number of layers increased, the performance increased. After 6 hidden layers the performance on the development set dropped. Hence we chose to fix the number of layers to 6. Thus, DNNs with 6 hidden layers were trained and the model performing better on the development set was chosen as the optimal model.

The DNN was trained on mini batches as the total amount of input features were quite high to be trained as a whole batch. The training was performed to minimize the mean squared error between the actual TVs and the estimated TVs. The Adam optimizer was used for optimizing the network parameters. Since the articulatory trajectories are low-pass in nature, the estimated TVs from the DNN were passed through a Kalman filter to remove additional noise [30]

3.4. Performance Measurement

The best performing DNN is the one that has the least mean squared error on the validation set. Since the dataset contained articulatory data from multiple speakers, correlation is a more appropriate performance measure than mean squared error. The Pearson Product Moment Correlation (PPMC) between the estimated TV and the corresponding ground-truth TV was computed as the performance evaluation metric.

4. Experiments and Results

The cortical feature space was initially reduced to a sub-space with 4, 5 and 7 number of scale, rate and frequency principal components respectively to obtain the best performing DNN model out of the models having 400, 512 and 800 nodes in each hidden layer for 6 hidden layers. To investigate the impact of the number of principal components being used in the vectorized features, we performed DNN training on 2 different sets of scale, rate, and frequency components of the input features: (4, 5, 7) and (4, 6, 8). Depending on this the size of the input layer was varied as 980 and 1344 nodes respectively. Dropout layers were used to reduce over-fitting.

Table 1: *Average correlations of estimated TVs for different DNN architectures, using (4, 5, 7), and (4, 6, 8) scale, rate and frequency coefficients for the cortical features*

# Hidden layer nodes	Cortical_980 (4,5,7)	Cortical_1344 (4,6,8)
400	0.635	0.685
512	0.645	0.684
800	0.632	0.694

The 6-layer DNN with 512 nodes in each hidden layer

yielded the best average correlation of 0.645 on the development set for 4,5 and 7 scale, rate and frequency coefficients.

According to Figure 2, the contribution from scale, rate and frequency coefficients of (4, 5, 7) respectively seems to represent the subspaces well. To get a more accurate representation of the input features, the rate and frequency components were increased to 6 and 8. The same experiments were repeated using 0.1 dropout for the input layer and 0.2 dropout for the subsequent layers. The results obtained on the development set are summarized in Table 1.

It can be seen that preserving more principal components from scale and frequency axes improves the correlation by 7.6%. The Table 2 summarizes and compares the average correlation values for each of the 6 TVs for different input features for speech inversion - Cortical_980, Cortical_1344, and MFCCs.

Table 2: Average Correlations for Estimated TVs (defined in Section 2)

Feature	Context frames	LA	LP	TBCL	TBCD	TTCL	TTCD	Average Correlation
Cortical_980 (4,5,7)	7	0.723	0.494	0.751	0.534	0.523	0.768	0.632
Cortical_980 (4,6,8)	7	0.758	0.516	0.789	0.589	0.576	0.822	0.675
MFCC	17	0.856	0.613	0.866	0.745	0.707	0.907	0.782

5. Conclusion

In this work we bring together the General Auditory Theory and the Motor Theory of speech perception. We develop a speech inversion system using the multi-resolution spectro-temporal ‘cortical features’ which are critical in the perception of sound in mammals. Experiments performed in the paper prove that these representations can estimate the articulatory features (TVs) with a correlation of 0.675 with ground-truth TVs.

The HOSVD dimensionality reduction preserves the slow varying components of speech which are ideal for estimating articulatory features. Since MFCCs directly correspond to capturing vocal tract variations and phonetic characteristics of speech signals, the correlations are much higher compared to the results obtained from our experiments in this paper. Moreover, the experiments with MFCC features had used a feature concatenation of 17 frames as against 7 frames for the cortical features. However the proper choice of rate and scale vectors has a significant impact to the correlation. The rates of 4, 8, 16 and 32 Hz capture the syllabic rate of the speech signals which corresponds to the tract variable trajectories. Hence with the appropriate parameter and feature dimensionality selection it can be seen that it is feasible to use the rich representation of the speech signals provided by cortical features to capture speech articulations.

In the future we plan to explore a joint optimization of the dimensionality reduction and the articulatory feature estimation so that the dimensionalities are optimally reduced for the speech inversion problem.

6. References

- [1] S. Aryal and R. Gutierrez-Osuna, “Accent conversion through cross-speaker articulatory synthesis,” in *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*. IEEE, may 2014, pp. 7694–7698.
- [2] M. Cavin, “The use of ultrasound biofeedback for improving English/t,” *Working Papers of the Linguistics Circle*, vol. 25, no. 1, pp. 32–41, 2015.
- [3] K. Kirchhoff, G. A. Fink, and G. Sagerer, “Combining acoustic and articulatory feature information for robust speech recognition,” *Speech Communication*, vol. 37, no. 3-4, pp. 303–319, jul 2002. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0167639301000206>
- [4] C. Qin and M. Á. Carreira-Perpiñán, “An empirical investigation of the nonuniqueness in the acoustic-to-articulatory mapping,” *Interspeech*, pp. 74–77, 2007.
- [5] B. S. Atal, J. J. Chang, M. V. Mathews, and J. W. Tukey, “Inversion of articulatory-to-acoustic transformation in the vocal tract by a computer-sorting technique,” *The Journal of the Acoustical Society of America*, vol. 63, no. 5, pp. 1535–1555, may 1978. [Online]. Available: <http://asa.scitation.org/doi/10.1121/1.381848>
- [6] S. Ouni and Y. Laprie, “Modeling the articulatory space using a hypercube codebook for acoustic-to-articulatory inversion,” *The Journal of the Acoustical Society of America*, vol. 118, no. 1, pp. 444–460, 2005.
- [7] Y. Laprie and B. Mathieu, “A variational approach for estimating vocal tract shapes from the speech signal,” *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP '98 (Cat. No.98CH36181)*, vol. 2, pp. 929–932, 1998. [Online]. Available: [papers://e7d065ae-9998-4287-8af0-c9fa85af8e96/Paper/p34600](https://doi.org/10.1109/ICASSP98.7287-8af0-c9fa85af8e96/Paper/p34600)
- [8] K. Richmond, “A trajectory mixture density network for the acoustic-articulatory inversion mapping,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, vol. 2, pp. 577–580. [Online]. Available: http://link.springer.com/10.1007/978-3-540-77347-4_23
- [9] S. Hiroya and M. Honda, “Estimation of Articulatory Movements From Speech Acoustics Using an HMM-Based Speech Production Model,” *IEEE Transactions on Speech and Audio Processing*, vol. 12, no. 2, pp. 175–185, mar 2004.
- [10] P. K. Ghosh and S. Narayanan, “A generalized smoothness criterion for acoustic-to-articulatory inversion,” *The Journal of the Acoustical Society of America*, vol. 128, no. 4, pp. 2162–72, oct 2010.
- [11] K. Kirchhoff, “Robust speech recognition using articulatory information,” 1999. [Online]. Available: <http://pub.uni-bielefeld.de/publication/2302713>
- [12] G. Sivaraman, V. Mitra, H. Nam, M. Tiede, and C. 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. [Online]. Available: <https://doi.org/10.1121/1.5116130>
- [13] N. Seneviratne, G. Sivaraman, V. Mitra, and C. Espy-Wilson, “Noise robust acoustic to articulatory speech inversion,” in *Proc. Interspeech 2018*, 2018, pp. 3137–3141. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1509>
- [14] N. Seneviratne, G. Sivaraman, and C. Espy-Wilson, “Multi-Corpus Acoustic-to-Articulatory Speech Inversion,” in *Proc. Interspeech 2019*, 2019, pp. 859–863. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-3168>
- [15] R. Santoro, M. Moerel, F. De Martino, R. Goebel, K. Ugurbil, E. Yacoub, and E. Formisano, “Encoding of Natural Sounds at Multiple Spectral and Temporal Resolutions in the Human Auditory Cortex,” *PLoS Computational Biology*, vol. 10, no. 1, p. e1003412, jan 2014. [Online]. Available: <http://dx.plos.org/10.1371/journal.pcbi.1003412>
- [16] R. L. Diehl, A. J. Lotto, and L. L. Holt, “Speech Perception,” *Annual Review of Psychology*, vol. 55, no. 1, pp. 149–179, feb 2004.
- [17] G. Fant, *Acoustic theory of speech production*. Walter de Gruyter, 1970, no. 2.
- [18] K. N. Stevens and S. E. Blumstein, “Invariant cues for place of articulation in stop consonants,” *The Journal of the Acoustical Society of America*, vol. 64, no. 5, pp. 1358–1368, 1978.
- [19] A. M. Liberman, F. S. Cooper, D. P. Shankweiler, and M. Studdert-Kennedy, “Perception of the speech code,” *Psychological review*, vol. 74, no. 6, 1967.

- [20] A. M. Liberman and I. G. Mattingly, "The motor theory of speech perception revised," *Cognition*, vol. 21, no. 1, pp. 1–36, 1985.
- [21] J. R. Westbury, "Speech Production Database User ' S Handbook," *IEEE Personal Communications - IEEE Pers. Commun.*, vol. 0, no. June, 1994.
- [22] R. 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] V. Mitra, H. Nam, C. Espy-Wilson, E. Saltzman, and L. 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, mar 2012. [Online]. Available: <http://asa.scitation.org/doi/10.1121/1.3682038>
- [24] T. Chi, P. Ru, and S. A. Shamma, "Multiresolution spectrotemporal analysis of complex sounds," *The Journal of the Acoustical Society of America*, vol. 118, no. 2, pp. 887–906, aug 2005. [Online]. Available: <http://asa.scitation.org/doi/10.1121/1.1945807>
- [25] P. Ru, "Multiscale Multirate Spectro-Temporal Auditory Model," Ph.D. dissertation, University of Maryland College Park, 2001.
- [26] K. Wang and S. A. Shamma, "Spectral Shape Analysis in the Central Auditory System," *IEEE Transactions on Speech and Audio Processing*, vol. 3, no. 5, pp. 382–395, 1995. [Online]. Available: <http://ieeexplore.ieee.org/document/466657/>
- [27] T. Chi, Y. Gao, M. C. Guyton, P. Ru, and S. Shamma, "Spectro-temporal modulation transfer functions and speech intelligibility," *The Journal of the Acoustical Society of America*, vol. 106, no. 5, pp. 2719–2732, nov 1999. [Online]. Available: <http://asa.scitation.org/doi/10.1121/1.428100>
- [28] N. Mesgarani, M. Slaney, and S. Shamma, "Discrimination of speech from nonspeech based on multiscale spectro-temporal Modulations," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 14, no. 3, pp. 920–930, may 2006. [Online]. Available: <http://ieeexplore.ieee.org/document/1621204/>
- [29] L. De Lathauwer, B. De Moor, and J. Vandewalle, "A Multilinear Singular Value Decomposition," *SIAM Journal on Matrix Analysis and Applications*, vol. 21, no. 4, pp. 1253–1278, jan 2000. [Online]. Available: <http://epubs.siam.org/doi/10.1137/S0895479896305696>
- [30] G. Sivaraman, "Articulatory representations to address acoustic variability in speech," Ph.D. dissertation, University of Maryland College Park, 2017. [Online]. Available: <https://drum.lib.umd.edu/handle/1903/20422>