

Audio Data Augmentation for Acoustic-to-articulatory Speech Inversion

1st Yashish M. Siriwardena

*Department of Electrical and Computer Engineering
University of Maryland College Park, MD, USA
yashish@umd.edu*

2nd Ahmed Adel Attia

*Department of Electrical and Computer Engineering
University of Maryland College Park, MD, USA
aadel@umd.edu*

3rd Ganesh Sivaraman

*Pindrop, GA, USA
ganesa90@gmail.com*

4th Carol Espy-Wilson

*Department of Electrical and Computer Engineering
University of Maryland College Park, MD, USA
espy@umd.edu*

Abstract—Data augmentation has proven to be a promising prospect in improving the performance of deep learning models by adding variability to training data. In previous work with developing a noise robust acoustic-to-articulatory speech inversion (SI) system, we have shown the importance of noise augmentation to improve the performance of speech inversion in ‘noisy’ speech conditions. In this work, we extend this idea of data augmentation to improve the SI systems on both the clean speech and noisy speech data by experimenting three data augmentation methods. We also propose a Bidirectional Gated Recurrent Neural Network as the speech inversion system instead of the previously used feed forward neural network. The inversion system uses mel-frequency cepstral coefficients (MFCCs) as the input acoustic features and six vocal tract-variables (TVs) as the output articulatory targets. The Performance of the system was measured by computing the correlation between estimated and actual TVs on the Wisconsin X-ray microbeam database. The proposed speech inversion system shows a 5% relative improvement in correlation over the baseline noise robust system for clean speech data. The pre-trained model, when adapted to each unseen speaker in the test set, improves the average correlation by another 6%.

Index Terms—data augmentation, noise-robust speech inversion, BiGRNN, vocal tract variables

I. INTRODUCTION

Acoustic-to-articulatory speech inversion (SI) is the task of retrieving the articulatory dynamics for a given speech signal [1]. The mapping from acoustics to articulation is an ill-posed problem which is known to be highly non-linear and non-unique [2]. However, over the recent years, the development of Speech Inversion (SI) systems has gained attention due to its potential in wide range of applications ranging from Automatic Speech Recognition (ASR) [3], speech synthesis [4], speech therapy [5] and mental health assessments [6], [7]. [6], [7]. Real articulatory data is obtained using techniques like X-ray microbeam [8], Electromagnetic Articulometry (EMA) [14] and real-time Magnetic Resonance Imaging (rt-MRI) [9]. These techniques are expensive, time consuming and require specialized equipment for observing articulatory movements

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

directly [1]. Hence, developing a speaker-independent SI system which can accurately estimate articulatory features for any unseen speaker can transform how speech research is conducted.

Advancements in deep neural networks (DNNs), especially in processing time series data to capture contextual information has propelled the development of SI systems to new heights. Bidirectional LSTMs (BiLSTMs) [10], CNN-BiLSTMs [11], Temporal Convolutional Networks (TCN) [12] and transformer models [13] have gained state-of-the-art results with multiple articulatory datasets [14]. In our previous work, we have reported SI results with the XRMB dataset [1] where a simple feed-forward neural network was trained with manually contextualized MFCCs as input features. In this work, we propose a Bidirectional Gated Recurrent Neural Network (BiGRNN) which outperforms the existing feed-forward neural network model on predicting TVs with the XRMB dataset. The BiGRNN model is light-weight, does not need any pre-contextualization of the input audio features or post-processing with low pass filters.

Majority of the SI systems are usually trained with a single corpus of data and have shown to perform poorly in cross-corpus [15] or speaker-independent experiments [11]. Difficulty collecting a larger corpus of data with many subjects and the differences in procedures of articulatory data collection (placement of sensors, EMA vs XRMB) have also made things complicated. To address the issue of data scarcity, the machine learning community has recently turned in the direction of data augmentation and generating synthetic data to train DNN models. The idea of training a DNN model on similar but different examples is known as data augmentation, which was initially proposed in [16] and later formalized by the Vicinal Risk Minimization (VRM) principle in [17]. Data augmentation has been widely utilized in improving the robustness of speech applications over the last few years (eg. automatic speech recognition (ASR) [18] and speech emotion recognition [19]). However, limited work has been done in the articulatory speech inversion domain, specially to improve the overall performance of SI systems with different, readily

available data augmentation techniques. In previous work with developing a noise robust articulatory SI system, we explored the idea of generating additive noise audio data [20]. The goal of that work was to develop a SI system which is robust against noisy acoustic data which is a common issue with rt-MRI articulatory datasets that contain significant amounts of MRI machine noise. Compared to the SI system trained entirely with clean speech data, the noise robust model did not perform well with the clean speech test set. The idea of developing a SI system which is robust against noisy speech data while also outperforming the current best performing model on clean speech data is the impetus for the current study.

II. DATASET DESCRIPTION

A. X-Ray Microbeam (XRMB) dataset

The original University of Wisconsin XRMB database [8] 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.

The anatomy of the speaker’s vocal tract defines the absolute positions of the articulators. 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 [21]. The TVs are based on articulatory phonology, a theoretical framework for speech production [22]. Using geometric transformations, the XRMB trajectories were converted to TV trajectories as outlined in [23]. 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).

III. AUDIO DATA AUGMENTATION

We conducted 3 types of audio data augmentations using the Audiomentation library¹.

A. Background Noise and Music

We added the noise and music data from MUSAN corpus [24] with the audio files in the XRMB dataset. The music files from the corpus were manually checked to remove any files with singing voices/human voices. We created two copies of the original speech data, one adding noise and the other




Fig. 1: Proposed BiGRNN model architecture

adding music. The gain of noise and music added are set to be proportional to the Root Mean Square (RMS) value of the input sound in the audio. The added noise/music can range from 5dB to 20dB SNR and the level is randomly chosen.

B. Gaussian Noise

Two copies of the original audio file are generated by adding Gaussian noise ranging from 5dB SNR to 20dB SNR. The noise SNR is sampled uniformly from the [5dB, 20dB] range when creating the noisy copy of the data.

C. Environmental Impulse Response (IR) functions

To add reverberation noise, we used the environmental IR functions from the MIT Acoustical Reverberation Scene Statistics Survey corpus [25]. Similar to the previous noise types, two noise copies of the original audio files were generated by convolving with randomly chosen IR functions from the corpus.

IV. SPEECH INVERSION SYSTEMS

A. Input Audio Features

All the audio files are first segmented to 2 second long segments and the shorter ones are zero padded at the end. Mel-Frequency Cepstral Coefficients (MFCCs) and Melspectrogram (MSPEC) features are then extracted as the acoustic input features for 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. Δ and $\Delta\Delta$'s for MFCCs were also computed to be used as an extended input feature set along with MFCCs. Both MFCCs (and Δ s) and MSPECs are utterance wise normalized (z-normalized) prior to model training. Table II shows the results for how well each audio feature type performed with the SI systems.

B. Proposed Bidirectional Gated RNN (BiGRNN) Model

Gated Recurrent Unit (GRU) was first proposed by Cho et al. [26] and has only 2 gates compared to 3 gates in a conventional LSTM unit, hence resulting in relatively smaller models and takes lesser time for training. In this paper, we propose a Bidirectional Gated RNN model as the speech inversion system. The model has 3 bidirectional layers of Gated Recurrent Units (GRUs) followed by two time distributed fully connected layers. Figure 1 shows the detailed model architecture used with the data augmented XRMB dataset. We specifically used a Masking layer at the input to avoid the affect of padded zeros for TV predictions. Dropout layers were used after every layer to minimize over-fitting.

C. Baseline BiLSTM and CNN-BiLSTM Models

To compare against the proposed BiGRNN model, we used a BiLSTM model inspired by the work of [10] and a CNN-BiLSTM model similar to that in [11] as the baseline models.

¹<https://github.com/iver56/audiomentations>

The BiLSTM model was trained with MFCCs and the CNN-BiLSTM model was trained with MSPECs.

The BiLSTM model has 3 bidirectional LSTM layers followed by two time distributed fully connected layers. The model has the same architecture as the proposed BiGRNN model except that it uses BiLSTM layers instead of the BiGRNN layers in the front-end of the model. The CNN-BiLSTM model consists of 5 CNN layers, whose outputs are then concatenated together and fed to 2 BiLSTM layers, followed by a final CNN layer.

All the models are evaluated with Mean Squared Error (MSE), Mean Absolute Error (MAE) and Pearson Product Moment Correlation (PPMC) scores computed between the estimated TVs and the corresponding ground-truth TVs.

D. Model Training

The input XRMB 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). None of the training, development and testing sets have overlapping speakers and hence all the models are trained in a ‘speaker-independent’ fashion. The split also ensured that around 80% of the total number of utterances were present in the training and the development and testing sets have a nearly equal number of utterances. This allocation was done in a completely random manner. When training the models with data augmentations as discussed in section III, all the augmented audio files were included in the same split as its original audio file to preserve the ‘speaker-independence’ and to not affect the original train-dev-test splitting ratios.

All the models were implemented with TensorFlow-Keras machine learning framework and trained with NVIDIA TITAN X GPUs. MSE, MAE and PPMC were experimented as loss functions to optimize the models. PPMC loss is defined as $1 - PPMC$ and was used as a custom loss to optimize the models. Table I shows the average correlation on the test set for the best performing models with the 3 types of loss functions. It can be seen that the BiLSTM and BiGRNN models converged better with MAE loss where as CNN-BiLSTM model worked better with the PPMC loss. One limitation with the PPMC loss is that it could predict TVs with an offset from the ground-truth which is expected (since it only captures variability more than the actual values). To address that we combined PPMC loss with the MAE loss where a weight ($0 < \alpha < 1$) was assigned to PPMC loss and $1 - \alpha$ was assigned to MAE loss. All the CNN-BiLSTM models are trained with the new weighted loss function which outperformed the MSE and MAE losses.

TABLE I: Comparison of Loss functions with each Model type

Model	Loss Fn.	MSE	MAE	PPMC	Weighted Loss ($\alpha = 0.8$)	Weighted Loss ($\alpha = 0.5$)
BiGRNN		0.7910	0.7959	0.7917	0.7928	0.7940
BiLSTM		0.7850	0.7870	0.7858	0.7862	0.7865
CNN-BiLSTM		0.7163	0.7214	0.7333	0.7256	0.7203

ADAM optimizer with a starting learning rate of 1e-3 and an exponential learning rate scheduler was used. The starting

learning rate was maintained up to 20 epochs and then decayed exponentially after each subsequent 5 epochs. The model was trained with an early stopping criteria (patience=5) monitoring the ‘validation loss’ on the development set. To choose the best starting ‘learning rate’ (LR), we did a grid search on [1e-3, 3e-4, 1e-4], whereas to choose the training batch size we did a similar grid search on [16,32,64,128]. Based on the validation loss, 1e-3 and 128 were chosen as the learning rate and batch size respectively.

V. EXPERIMENTS AND RESULTS

A. Performance comparison of different model architectures with different audio features

Previous work with developing SI systems have shown MFCCs to be superior over MSPECs and Perceptual Linear Predictions (PLPs) as acoustic features [1]. Our results with different audio features as shown in Table II are consistent with the previous studies which suggests MFCCs can be effective specially with Bidirectional RNN based models. We also show that adding Δ and $\Delta\Delta$ features derived from MFCCs does not necessarily improve the performance. Table II also presents the results of how different model architectures when trained with ‘clean speech’ data performs on unseen speaker data. The BiGRNN model had a clear advantage over the two baseline models (BiLSTM and CNN-BiLSTM) and hence was selected as the proposed model architecture for the SI system. It should also be noted, that in this experiment, all the models were trained with only clean speech data (with no data augmentation) to disentangle any affects of data augmentation to choose the best DNN architecture.

TABLE II: PPMC scores for models trained with clean speech when tested with clean speech test set

Model	Audio features	LA	LP	TBCL	TBCD	TTCL	TTCD	Average
BiGRNN	MFCC	0.8801	0.6200	0.8580	0.7382	0.6922	0.9206	0.7848
BiLSTM	MFCC	0.8742	0.6236	0.8535	0.7189	0.6792	0.9142	0.7773
BiGRNN	MFCC + Δ + $\Delta\Delta$	0.8708	0.6256	0.8565	0.7167	0.7020	0.9125	0.7807
BiLSTM	MFCC + Δ + $\Delta\Delta$	0.8480	0.6112	0.8347	0.7055	0.6667	0.8975	0.7606
CNN-BiLSTM	Melspectrogram	0.8285	0.5651	0.8028	0.6827	0.6193	0.8551	0.7256

B. Effectiveness of different types of audio data augmentations

One of the key contributions of this work is the use of audio data augmentations to improve the standard SI task. But to choose the right data augmentation type, the best performing BiGRNN SI system was independently trained with speech data augmented with the three proposed data augmentation techniques in section III. The individual SI systems were then independently tested with the clean speech test split. Figure 2 shows the PPMC scores for each TV, and the resulting average score from the BiGRNN model trained and evaluated in this fashion. Results suggest that adding background noise and music from MUSAN dataset performs slightly better compared to adding reverberation noise from IR functions or random Gaussian noise.

A second experiment was also done to evaluate each SI system trained with one augmentation type, tested with the data augmented with the other two types. As expected, from Table III, it is clear that the models trained with a certain augmentation type tends to perform well with the data augmented in the same fashion. Additionally, an average PPMC




Fig. 2: PPMC scores for each TV and the average score from the proposed BiGRNN model on clean speech test set

scores across different augmentation splits was also computed as shown in Table III. The average PPMC scores further support the fact that adding background noise and music is better choice compared to the other two data augmentation types considered.

TABLE III: PPMC scores for models trained with one type of data augmentation, tested with the other types

Train Aug.	Test Aug. Background Noise & Music	Gaussian Noise	Room IR	Average
Background Noise & Music	0.7830	0.7735	0.7572	0.7712
Gaussian Noise	0.7726	0.7843	0.7395	0.7655
Room IR	0.7624	0.7337	0.7797	0.7586

C. Performance of proposed data augmentations

Since data augmentation with background noise and music resulted in the best SI models, we trained all the subsequent models with data augmented in that fashion. Figure 3 shows the average PPMC score across the 6 TVs for the proposed BiGRNN, BiLSTM and CNN-BiLSTM models when trained with clean speech only (clean-train), and clean speech + augmented data (augment-train). The two SI systems for each model are tested with clean only, clean+augmented and augmented only test sets to evaluate the robustness of the models not only for noisy/augmented speech, but also for clean speech data. The best performing SI model is reported with the BiGRNN model, trained with augmented data, and it predicts TVs with an average PPMC score of 0.7959 on clean speech data. For the context, the previous noise-robust model in [20] was only able to achieve a best PPMC score of 0.741 on the clean test set. Therefore, the new SI system gains a 5% relative improvement over the previous noise-robust SI system on estimating TVs for clean speech data. This elucidates the fact that, data augmentation can be used as a promising technique not only for making SI systems noise-robust, but also for improving the general SI task on clean speech data.

D. Model adaptation for unseen-target speakers

In previous work with SI systems, it has been shown that model adaptation with a generalized (pre-trained) model trained with a larger number of subjects can perform better than a speaker-dependent model trained only with target speaker’s data [10]. Based on that observation, we performed model adaptation for speakers in the test set to see how much

of an improvement the models can gain when an already pre-trained model is further trained with a portion of the subject’s data.

We take the best performing BiGRNN model pre-trained with augmented data and further train the model with each speaker’s data in the test set to create individual speaker adapted models. The pre-trained model weights are only used for initialization and the new training was carried out with a smaller batch size (=4), reduced starting LR (=1e-4) and a quick decay with LR scheduler (every 2 epochs). The early stopping patience was increased to 30 to account for slight fluctuations in the validation loss due to smaller amount of training data. Here we used 80 % of all the subject’s data (average 12 mins of speech) for training and used 10 % each for validation and testing. Table IV shows the PPMC results of BiGRNN models when tested with same test split of subject’s data before and after model adaptation. Turns out for the 5 speakers in the test set, an average PPMC score of 0.8506 can be achieved when the speaker-adapted models are tested with held-out data from the same subject. The pre-trained model without any adaptation, tested on the same splits of the target subjects can only achieve a PPMC score of 0.7942.

Figure 4 shows ground-truth and predicted TVs, LA, TBCD, and TTCD for an example utterance of JW31 subject. The TVs are estimated by the pre-trained and the speaker-adapted SI systems. The effect of speaker adaptation is clearly evident with the estimated TV trajectories, where the predicted TVs from the speaker adapted model looks significantly better than that estimated by the generalized model. However, it should also be noted, that for certain speakers (e.g. JW61) this improvement is not clearly evident. Hence, further work needs to be done to understand what speaker-specific characteristics are captured by these speaker-adapted SI systems and to devise effective modifications to infuse speaker characteristics to improve speaker-adapted SI systems.

TABLE IV: PPMC score for TV predictions, before and after model adaptation

Subject	Before Model Adaptation	After Model Adaptation
JW31	0.8071	0.9106
JW39	0.7907	0.8836
JW18	0.7936	0.8633
JW33	0.8334	0.8490
JW61	0.7463	0.7466

VI. CONCLUSION AND FUTURE WORK

In this work, we propose a BiGRNN model architecture to perform acoustic-to-articulatory speech inversion on the XRMB dataset. We show that audio augmentations can noticeably improve the SI task to perform well in both noisy and clean speech data. We also compare our proposed model with two baseline model architectures widely used in the SI domain. Finally, with speaker adaptation, we show that the performance of the BiGRNN model can be further improved for unseen speakers. As future work, we are hoping to conduct experiments with multiple articulatory corpora to explore how




Fig. 3: Performance of each model with and without augmented data on clean, augmented and clean+augmented test sets




Fig. 4: TV plots for the clean utterance ‘combined are the ingredients in a large bowl’ estimated by the SI systems before and after adaptation on the JW31 subject

well data augmentation with multiple datasets can help in better generalizability of SI systems.

REFERENCES

- [1] 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>
- [2] C. Qin and M. A. Carreira-Perpiñán, “An empirical investigation of the nonuniqueness in the acoustic-to-articulatory mapping,” *Interspeech*, pp. 74–77, 2007.
- [3] V. Mitra, C. Y. Espy-Wilson, E. Saltzman, and L. Goldstein, “Articulatory information for noise robust speech recognition,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 1913–1924, 2011.
- [4] Z.-H. Ling, K. Richmond, and J. Yamagishi, “Articulatory control of hmm-based parametric speech synthesis using featurespace-switched multiple regression,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 1, pp. 207–219, 2013.
- [5] S. Fagel and K. Madany, “A 3-d virtual head as a tool for speech therapy for children,” in *INTERSPEECH*, 2008.
- [6] C. Espy-Wilson, A. C. Lammert, N. Seneviratne, and T. F. Quatieri, “Assessing Neuromotor Coordination in Depression Using Inverted Vocal Tract Variables,” in *Proc. Interspeech 2019*, 2019, pp. 1448–1452.
- [7] Y. M. Siriwardena, C. Espy-Wilson, C. Kitchen, and D. L. Kelly, *Multi-modal Approach for Assessing Neuromotor Coordination in Schizophrenia Using Convolutional Neural Networks*. New York, NY, USA: Association for Computing Machinery, 2021, p. 768–772. [Online]. Available: <https://doi.org/10.1145/3462244.3479967>
- [8] J. R. Westbury, “Speech Production Database User’s Handbook,” *IEEE Personal Communications - IEEE Pers. Commun.*, vol. 0, no. June, 1994.
- [9] S. Narayanan, K. Nayak, S. Lee, A. Sethy, and D. 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.
- [10] A. Illa and P. K. Ghosh, “Low Resource Acoustic-toarticulatory Inversion Using Bi-directional Long Short Term Memory,” in *Proc. Interspeech 2018*, 2018, pp. 3122–3126.
- [11] A. S. Shahrehabaki, S. M. Siniscalchi, G. Salvi, and T. Svendsen, “Sequence-to-Sequence Articulatory Inversion Through Time Convolution of Sub-Band Frequency Signals,” in *Proc. Interspeech 2020*, 2020, pp. 2882–2886.
- [12] A. S. Shahrehabaki, S. M. Siniscalchi, and T. Svendsen, “Raw Speech-to-Articulatory Inversion by Temporal Filtering and Decimation,” in *Proc. Interspeech 2021*, 2021, pp. 1184–1188.
- [13] S. Udupa, A. Roy, A. Singh, A. Illa, and P. K. Ghosh, “Estimating Articulatory Movements in Speech Production with Transformer Networks,” in *Proc. Interspeech 2021*, 2021, pp. 1154–1158.
- [14] M. Tiede, C. Y. Espy-Wilson, D. Goldenberg, V. Mitra, H. Nam, and G. 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.
- [15] N. Seneviratne, G. Sivaraman, and C. Espy-Wilson, “MultiCorpus Acoustic-to-Articulatory Speech Inversion,” in *Proc. Interspeech 2019*, 2019, pp. 859–863.
- [16] P. Y. Simard, Y. A. LeCun, J. S. Denker, and B. Victorri, *Transformation Invariance in Pattern Recognition – Tangent Distance and Tangent Propagation*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 235–269.
- [17] O. Chapelle, J. Weston, L. Bottou, and V. Vapnik, “Vicinal risk minimization,” in *Advances in Neural Information Processing Systems*, T. Leen, T. Dietterich, and V. Tresp, Eds., vol. 13. MIT Press, 2000.
- [18] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, “Audio augmentation for speech recognition,” in *INTERSPEECH*, 2015.
- [19] R. Pappagari, J. Villalba, P. Zelasko, L. Moro-Velazquez, and N. Dehak, “Copypaste: An augmentation method for speech emotion recognition,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6324–6328.
- [20] 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.
- [21] 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.
- [22] C. P. Browman and L. Goldstein, “Articulatory Phonology : An Overview *,” *Phonetica*, vol. 49, pp. 155–180, 1992.
- [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, 2012.
- [24] D. Snyder, G. Chen, and D. Povey, “MUSAN: A Music, Speech, and Noise Corpus,” 2015, arXiv:1510.08484v1.
- [25] J. Traer and J. H. McDermott, “Statistics of natural reverberation enable perceptual separation of sound and space,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 48, pp. E7856–E7865, 2016.
- [26] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the properties of neural machine translation: Encoder–decoder approaches,” in *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Doha, Qatar: Association for Computational Linguistics, Oct. 2014, pp. 103–111.