



Aphasic Speech Recognition using a Mixture of Speech Intelligibility Experts

Matthew Perez, Zakaria Aldeneh, Emily Mower Provost

University of Michigan, Ann Arbor

mkperez@umich.edu, aldeneh@umich.edu, emkprovost@umich.edu

Abstract

Robust speech recognition is a key prerequisite for semantic feature extraction in automatic aphasic speech analysis. However, standard one-size-fits-all automatic speech recognition models perform poorly when applied to aphasic speech. One reason for this is the wide range of speech intelligibility due to different levels of severity (i.e., higher severity lends itself to less intelligible speech). To address this, we propose a novel acoustic model based on a mixture of experts (MoE), which handles the varying intelligibility stages present in aphasic speech by explicitly defining severity-based experts. At test time, the contribution of each expert is decided by estimating speech intelligibility with a speech intelligibility detector (SID). We show that our proposed approach significantly reduces phone error rates across all severity stages in aphasic speech compared to a baseline approach that does not incorporate severity information into the modeling process.

Index Terms: disordered speech recognition, aphasia, speech intelligibility, mixture of experts, automatic speech recognition

1. Introduction

Aphasia is an acquired language disorder that manifests itself in speech [1, 2]. In the U.S., approximately two million people are living with aphasia and more than 180,000 acquire it every year due to brain injury, most commonly from a stroke [3]. Aphasic speech can be difficult to understand, and as a result, persons with aphasia (PWAs) are typically characterized as having low speech intelligibility. Previous studies have shown that speech-language therapy activities have a positive effect on the communication abilities of PWAs [4, 5]. However, current speech-language analyses and speech therapy activities require in-clinic visits with a trained speech-language pathologist, which is both costly and time-consuming. Speech-based technology provides an attractive avenue for assisting medical professionals due to its low cost, high accessibility, and promising results for automating processes that relate to the assistance and analysis of those with disordered speech [6–11]. Many of these automatic analyses rely on accurate speech transcriptions, which motivates the need for better automatic speech recognition (ASR).

There are several challenges associated with building robust ASR systems for aphasic speech, which include low speech intelligibility [12], high (inter- and intra-) speaker variability [13, 14], and relatively limited data [15]. Some artifacts of low speech intelligibility include halting speech, use of jargon, as well as various phone-level and word-level substitutions (paraphasias). These speech impairments can also be compounded by co-occurring motor control disorders such as apraxia of speech and dysarthria [16]. Additionally, in regards to machine learning efforts, it's difficult to apply large, monolithic networks that traditionally take advantage of big data due to the low-resource domain of aphasic speech.

Prior works have focused on improving ASR for disordered speech by addressing the issue of speaker variability through

speaker-selective model adaptation [15] and speaker embeddings [10]. In modeling speaker variability, it is likely that these models also capture speech intelligibility, as it is an underlying feature of speaker variability. Therefore, we believe that the success of these approaches demonstrates the potential for modeling speech intelligibility for disordered speech recognition.

In this work, we present a novel approach that explicitly models speech intelligibility in the acoustic model of an ASR system to provide more robust aphasic speech recognition. To accomplish this, we use an acoustic model architecture designed to handle the wide range of speech intelligibility present in aphasic speech. Specifically, we use a Mixture of Experts (MoE), deep neural network (DNN) acoustic model where each expert in the model focuses on specific severity classes defined using the speaker Aphasia Quotient (AQ) severity metric [17]. We introduce a *speech intelligibility detector* (SID), trained to detect severity levels of a given speech frame, to automatically estimate the contribution of each expert in our proposed acoustic model. Our results demonstrate that the proposed MoE acoustic model improves phone recognition performance over a traditional one-size-fits-all acoustic model across all speaker severity levels.

2. Related Works

Disordered speech recognition is an active area of study due to its important role in automatic speech-based analyses and potential impact on remote health monitoring applications. The disordered speech recognition domain shares many challenges with the accented speech recognition domain. Several approaches such as joint modeling, multitask-learning, and MoEs have been proposed in the field of accented speech recognition to address these challenges [18–20].

Le and Mower Provost improved aphasic speech recognition performance by adapting the acoustic model through appending input acoustic features with fixed-length speaker identity vectors (i-vectors) [10]. i-vectors reduce speaker variability in an acoustic model by providing information about general speaker characteristics (e.g., pronunciation patterns) [21]. Christensen et. al. showed that speaker selective training improves isolated word recognition on disordered speech [15]. The authors applied maximum a posteriori (MAP) adaptation to various HMM-GMM acoustic models with a select pool of speakers. They found that speaker-independent models trained on non-dysarthric speech were inferior to models trained on dysarthric speech for the task of disordered speech recognition, highlighting the importance of acoustic similarity in disordered speech recognition.

Findings from previous research suggest that careful consideration should be taken when training acoustic models for disordered speech applications as a high variation in speech intelligibility exists among speakers. One approach to address the high variation present in disordered speech is to train several independent acoustic models, one for each intelligibility group.

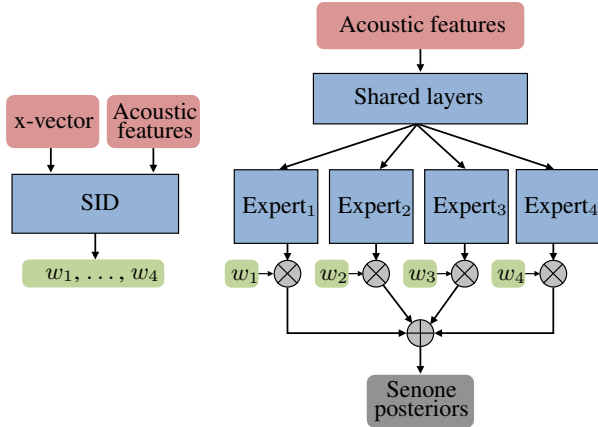


Figure 1: The proposed speech intelligibility identifier (SID) and mixture of experts acoustic model. The contributions of the individual experts is determined by the outputs from the SID.

However, the main challenge with training independent models is the limited amount of data available for each group. The proposed MoE acoustic model mitigates this challenge by allowing us to train a single network that explicitly decomposes the acoustic model into several *expert* acoustic models, where each expert is focused on a specific range of speech intelligibility. MoE based approaches in acoustic modeling have proven to be successful in accented speech recognition tasks by allowing acoustic models to handle several accents simultaneously [20]. To the best of our knowledge, we are the first to explore utilizing speech intelligibility in an acoustic model for aphasic speech recognition.

3. Data

AphasiaBank is a large audiovisual database composed of multiple sub-databases, which contain interactions between PWAs and clinicians [22]. Sub-databases contain a mix of speakers with varying severity and aphasia type (i.e., Broca’s, Conduction, etc.). Some sub-databases also include healthy speakers who followed a similar elicitation protocol. The elicitation protocol consists of free speech (i.e., personal story, describing current day), semi-structured speech (i.e., picture description tasks), and procedural discourse (i.e. “How would you make a peanut butter and jelly sandwich?”). In this work, we focus on English only sub-databases containing at least four speakers.

The final dataset we use contains 25 sub-databases with 106,509 utterances and over 95 hours of speech. There are 537 speakers (237 healthy, 138 mild-aphasia, 115 moderate-aphasia, 46 severe-aphasia). Of these speakers, 241 are male and 210 are female, with the average age being 62.5 ± 14.5 years. Our metric for speech intelligibility is the WAB-R Aphasia Quotient (AQ), which is a measurement for aphasic severity [17]. AQ scores ≥ 75 are considered mild, scores of 75-50 are considered moderate, and scores < 50 are considered severe.

We randomly split the dataset into training, validation, and test partitions using a 70/5/25 rule. Our test set is formed by randomly selecting 25% of the data in each sub-database. In addition, our splits are speaker-independent and severity-stratified to ensure similar distributions of severity across the three partitions. All recordings were down-sampled to 16kHz in this work.

Table 1: Data sharing protocols used when training the MoE acoustic models. M_i represents the i -th expert in the model. h , mi , mo , and s represent utterances from speakers that are Healthy, Mild, Moderate, and Severe, respectively.

Setup	M_0	M_1	M_2	M_3
Solo	h	mi	mo	se
Solo+Healthy	h	$h \cup mi$	$h \cup mo$	$h \cup se$
Solo+Neighbor	h	$h \cup mi$	$mi \cup mo$	$mo \cup se$

4. Method

The architecture for the proposed MoE acoustic model is shown in Figure 1. The inputs to the model are acoustic features and the outputs are context-dependent triphone state probabilities (senones). The acoustic model consists of two parts: (1) shared initial layers used for feature extraction and (2) several expert networks each focusing on a pre-defined severity group. We fix the number of experts to four (healthy, mild, moderate, and severe) following the speech intelligibility levels defined in Section 3. The shared layers of the model are designed to capture broad phonetic representations applied across all severities, while the experts are designed to reduce speech intelligibility variation within their respective networks. All hidden layers are fully-connected, consist of 1024 nodes, and use a ReLU activation function. The shared layers of the model consist of four hidden layers and each expert is composed of two hidden layers and the final output layer applies a softmax function. Our baseline network is a fully connected DNN with 6 hidden layers.

In addition to the acoustic model, we also propose a speech intelligibility detector (SID) which predicts speech intelligibility levels, as quantified by the AQ severity metric, given the acoustic features and utterance-level speaker embeddings. The outputs from the SID are used as weights for computing the contribution of each expert in our acoustic model. We use speaker embeddings extracted from a pre-trained x-vector system [23] and apply principal component analysis (PCA) to reduce the dimensionality of these x-vectors from 512 to 32 to avoid overfitting. x-vectors have previously been used to estimate speech intelligibility [24] due to their ability to encode many properties about a speaker (i.e. speaking style, speaking rate, etc.) [25,26]. The SID consists of two hidden layers and a final output layer which applies a softmax function.

The output of the acoustic model at time frame x is determined using the following equation:

$$y(x) = \sum_{i=0}^N w_i(x) M_i(x) \quad (1)$$

where $y(x)$ are the final senone probabilities, $w_i(x)$ is the contribution of expert i as computed by the SID, M_i are the senone probabilities from expert i , and N is the total number of experts in the model.

5. Experimental Setup

We use the Kaldi toolkit [28] for feature extraction and HMM-GMM training and use the pytorch-kaldi toolkit [29] for neural network training. The feature extraction pipeline involves the extraction of 13-dimensional mel frequency cepstral coefficient (MFCC) features, with cepstral mean and variance normalization (CMVN) spliced across 7 frames (± 3). Frames are created using a sliding window of size 25 ms with a 10 ms shift rate. We

Table 2: Model performance (in PER) based on severity classes. Relative PER improvement is shown in parenthesis and an * is used to indicate a PoI of >99% over the baseline, using the method proposed in [27]. Best performance for each setup is in **bold**.

Model	Overall	Mild	Moderate	Severe
Baseline	37.96	34.19	42.45	66.56
<i>Solo</i>	40.29 (-5.8)	36.21 (-5.6)	44.79 (-5.2)	77.38 (-13.9)
<i>Solo+Healthy</i>	38.35 (-1.0)	34.54 (-1.0)	42.77 (-0.7)	70.03 (-4.9)
<i>Solo+Neighbor</i>	37.03 (+2.5)*	33.37 (+2.4)*	41.69 (+1.8)*	61.41 (+7.7)*

apply Linear Discriminant Analysis (LDA) to reduce the dimension size to 40 and use Maximum Likelihood Linear Transformation (MLLT) to further decorrelate the features [30]. Lastly, we apply feature-space Maximum Likelihood Linear Regression (fMLLR) [31] as a means of speaker adaptation and use the resulting output as acoustic features for HMM-DNN training. The final inputs to the DNN acoustic models are 40-dimensional fMLLR acoustic features with a context window of 11 (± 5) frames, for a total feature size of 440.

We train our networks with an SGD optimizer using an initial learning rate of 0.01. The learning rate is halved after each epoch when the performance on the validation set stagnates. We use early stopping to reduce overfitting. For evaluation, we use Phone Error Rate (PER), following the same approach as [10]. To assess the statistical significance of PER improvement between our models, we use a bootstrap estimation test provided by the Kaldi `compute-wer-bootstrap` tool [27]. This method computes *probability of improvement* (PoI), where PER is computed on 10^4 sampled utterances for a given model in order to estimate the spread of PER around its mean. The difference in PER is then computed for each sample between the baseline and the proposed model being evaluated, providing an estimate of the probability of improvement.

6. Experiments and Results

6.1. MoE with Oracle Speech Intelligibility Detector

In our initial set of experiments, we assume that we have access to an *oracle* SID, which uses the speaker’s known AQ to select the proper expert at test time. The goal of this section is two-fold: (1) explore the relationship between data assignment among experts and acoustic model performance during training and (2) demonstrate the effectiveness of incorporating information about speech intelligibility into an MoE acoustic model. We experiment with several models to study the effect of data assignment among experts.

The *baseline* model that we start with has the same architecture as the model described in Section 4 with the exception of only having one expert. The baseline is a one-size-fits-all model that is typically used for general-purpose ASR applications. Next, we introduce a series of data assignment methods for training the MoE, outlined in Table 1.

The first approach, *Solo*, is where each expert in the MoE acoustic model only sees data samples from one severity group. The *Solo* model can be thought of as the other extreme to our baseline model; whereas the baseline model pools all data regardless of its severity group, the *Solo* model splits data among experts based on their respective severity group.

One challenge with the *Solo* approach, however, is that each expert in the model will see a limited amount of data. To compensate for this challenge, we introduce the *Solo + Healthy* approach, which augments the data seen by each ex-

pert with healthy speech samples. While this might improve the performance over the *Solo* approach, we hypothesize that large variation in speech intelligibility present among healthy and aphasic speakers might hinder model performance.

We finally introduce the *Solo+Neighbor* approach, which is used to augment the data seen by each expert while addressing the data discrepancy challenges that could arise from using the *Solo + Healthy* approach. In the *Solo + Neighbor* approach, we augment the data seen by each expert by incorporating data that is similar in terms of speech intelligibility. Specifically, each expert in our *Solo + Neighbor* acoustic model sees both its respective data and data from the neighboring expert (i.e., the expert focusing on lower severity class).

6.1.1. Results

Table 2 shows the model PER performance across all severity classes as well as the relative improvement and PoI over the baseline. We find that a MoE acoustic model, combined with an oracle SID and proper data assignment training (*Solo + Neighbor*), results in significantly lower PERs across all severity levels when compared to PERs obtained using a standard one-size-fits-all acoustic model.

First, we find that the *Solo* approach yields higher PERs compared to those obtained from the baseline. We believe this drop in performance is due to a lack of training data in each expert as the restrictive nature of *Solo* only exacerbates a data scarcity problem. As we move to a less restrictive data assignment method (*Solo + Healthy*), we can see improved performance over the *Solo* approach as healthy data is given to each expert. However, in comparison to the baseline, *Solo + Healthy* still performs worse, which we believe is due to the dissimilarity (in speech intelligibility) between healthy samples and the original *Solo* samples in each expert. This is confirmed by the performance of *Solo + Neighbor*, which shows that less restrictive data assignment that attempts to preserve speech intelligibility similarity yields the best results. With *Solo + Neighbor*, we see an overall 2.5% relative improvement over the baseline with 2.4%, 1.8%, and 7.7% relative improvements for mild, moderate, and severe forms of aphasia, respectively.

The results show that when training an MoE based on speech intelligibility, being overly restrictive with expert training can compound the effects of the data scarcity problem. We show that this can be alleviated through non-mutually exclusive data assignment and that preserving speech intelligibility consistency in each expert when considering pooling strategies (shown by *Solo + Neighbor*) leads to better performance over a one-size-fits-all model. With these considerations, we are able to show that a MoE acoustic model trained using the *Solo + Neighbor* approach improves phone recognition performance over a baseline when supported by an oracle SID that estimates speaker severity classes.

Table 3: Model performance (in PER) when using a trained SID at test time. All models are trained following the Solo + Neighbor approach. * indicates a PoI of >99% over the baseline, using the method proposed in [27]. Best performance for each setup is in **bold**.

Model	Overall	Mild	Moderate	Severe
Baseline	37.96	34.19	42.45	66.56
SID _{frame}	36.98 (+2.6)*	33.20 (+2.9)*	41.67 (+2.0)*	62.71 (+5.8)*
SID _{utt}	36.87 (+2.9)*	33.07 (+3.3)*	41.64 (+2.0)*	62.86 (+5.6)*

6.2. MoE with Automatic Speech Intelligibility Detector

Our previous experiments assumed that we had access to an oracle SID that perfectly predicts speech intelligibility. In this section, the contributions of each expert in our acoustic model are guided by the predictions made by a neural SID model that is discriminatively trained to predict severity classes given the same features used by the acoustic model in addition to speaker embeddings extracted by a pre-trained x-vector model.

We use the *Solo + Neighbor* data assignment approach when training our MoE, as it provided significant performance improvements over other setups in Section 6.1.1. We experiment with two variants of the SID in order to study the effect that averaging speech intelligibility has on ASR performance. More specifically, we investigate averaging SID predictions at the frame-level and at the utterance-level. Utterance-level SID predictions force the contributions of the experts to be consistent across all frames from a given utterance, while frame-level SID predictions can result in different expert contributions from one frame to next. We refer to the frame-level and utterance-level SIDs as *SID_{frame}* and *SID_{utt}*, respectively.

6.2.1. Results

In section 6.1, we assumed that the model had access to an oracle classifier, which was able to correctly identify the severity class of the speaker. Here, we train a SID using fMLLR and x-vector features to predict severity class at the frame-level. The performance of the SID is shown in Figure 2 where we see confusion between neighboring class severities. Although the SID is imperfect, the confusion trend between adjacent classes indicates that the SID is learning some aspects related to speech intelligibility. We believe that the poor SID performance is partly due to the use of AQ scores, which include other modalities beyond acoustic speech (i.e. language usage, motor control tests, family history, etc.). In addition to this, we believe another contributing factor is that these AQ labels are speaker-level speech intelligibility scores, which assumes that speech intelligibility manifests itself at the speaker-level when, in fact, speech intelligibility may vary on a much smaller time scale. Without access to smaller time-scale intelligibility scores, we investigate the impact that averaging the SID output at the frame-level and utterance-level have on MoE performance.

For all experiments in this section, the neural SID outputs, at test time, are used as soft weights (probabilities) for merging the senone predictions of all experts appropriately for a given test sample. Table 3 shows the PER, relative performance improvement, and PoI over the baseline, across severity classes.

SID_{frame} achieves improved results over the baseline model, across all severities, however, *SID_{utt}* yields slightly better results with 2.9% relative improvement overall with 3.3%, 2.0%, and 5.6% relative improvements for mild, moderate, and severe respectively. This shows that a trained neural SID can be used to improve phone recognition in an MoE acoustic model over a traditional one-size-fits-all model. It is also

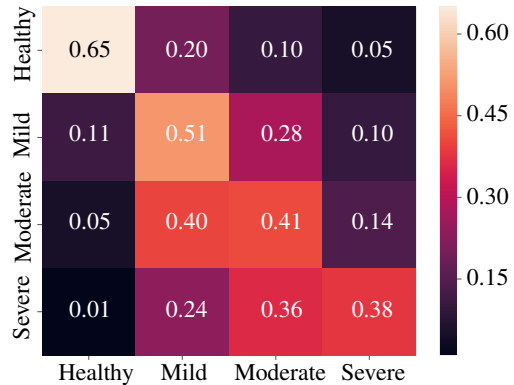


Figure 2: Confusion Matrix of frame-level SID performance

interesting to note that using an automatic SID outperformed using an oracle SID in the MoE framework. We believe this may indicate that speaker-level speech intelligibility scores may not be as accurate as utterance-level or frame-level intelligibility scores. Overall, these results show that despite the imperfect performance of the SID at predicting speaker-level speech intelligibility, the MoE still achieves performance improvement over a baseline model.

7. Conclusion

In this work, we first show that modeling speech intelligibility in a MoE acoustic model improves phone recognition across all severities with the use of an oracle SID. We then demonstrate the importance of data assignment when training the MoE model. As we move away from the use of an oracle SID, we show how a neural SID can be trained using fMLLR and x-vector features to estimate severity classes. Lastly, we show that the outputs of the SID can be used to merge expert senone predictions at test time for improved MoE performance over a one-size-fits-all acoustic model.

In section 6.2.1 we suggest that a potential limitation of this work is in utilizing speaker-level speech intelligibility scores. In future work, we plan to explore smaller time-scale metrics for estimating speech intelligibility and how this can be used to improve the SID and ultimately our MoE acoustic model.

8. Acknowledgments

This material is based in part upon work supported by the Toyota Research Institute (“TRI”) and by the National Science Foundation Graduate Research Fellowship Program. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF, TRI, or any other Toyota entity.

9. References

- [1] M. Danly and B. Shapiro, "Speech prosody in broca's aphasia," *Brain and language*, vol. 16, no. 2, pp. 171–190, 1982.
- [2] S. Ash, C. McMillan, D. Gunawardena, B. Avants, B. Morgan, A. Khan, P. Moore, J. Gee, and M. Grossman, "Speech errors in progressive non-fluent aphasia," *Brain and language*, vol. 113, no. 1, pp. 13–20, 2010.
- [3] N. A. Association, "Aphasia," 2020.
- [4] A. Basso, "Prognostic factors in aphasia," *Aphasiology*, vol. 6, no. 4, pp. 337–348, 1992.
- [5] S. K. Bhogal, R. Teasell, and M. Speechley, "Intensity of aphasia therapy, impact on recovery," in *Database of Abstracts of Reviews of Effects (DARE): Quality-assessed Reviews [Internet]*, Centre for Reviews and Dissemination (UK), 2003.
- [6] H. Christensen, I. Casanueva, S. Cunningham, P. Green, and T. Hain, "Automatic selection of speakers for improved acoustic modelling: Recognition of disordered speech with sparse data," in *2014 IEEE Spoken Language Technology Workshop (SLT)*, pp. 254–259, IEEE, 2014.
- [7] D. Le, K. Licata, C. Persad, and E. M. Provost, "Automatic assessment of speech intelligibility for individuals with aphasia," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 11, pp. 2187–2199, 2016.
- [8] D. Le, K. Licata, and E. M. Provost, "Automatic quantitative analysis of spontaneous aphasic speech," *Speech Communication*, vol. 100, pp. 1–12, 2018.
- [9] K. C. Fraser, F. Rudzicz, N. Graham, and E. Rochon, "Automatic speech recognition in the diagnosis of primary progressive aphasia," in *Proceedings of the fourth workshop on speech and language processing for assistive technologies*, pp. 47–54, 2013.
- [10] D. Le and E. M. Provost, "Improving automatic recognition of aphasic speech with aphasiabank," in *Interspeech*, pp. 2681–2685, 2016.
- [11] M. Perez, W. Jin, D. Le, N. Carozzi, P. Dayalu, A. Roberts, and E. M. Provost, "Classification of huntington disease using acoustic and lexical features," in *Interspeech*, pp. 1898–1902, 2018.
- [12] K. T. Mengistu and F. Rudzicz, "Comparing humans and automatic speech recognition systems in recognizing dysarthric speech," in *Canadian Conference on Artificial Intelligence*, pp. 291–300, Springer, 2011.
- [13] M. S. Hawley, P. Enderby, P. Green, S. Cunningham, S. Brownsell, J. Carmichael, M. Parker, A. Hatzis, P. O'Neill, and R. Palmer, "A speech-controlled environmental control system for people with severe dysarthria," *Medical Engineering & Physics*, vol. 29, no. 5, pp. 586–593, 2007.
- [14] M. B. Mustafa, F. Rosdi, S. S. Salim, and M. U. Mughal, "Exploring the influence of general and specific factors on the recognition accuracy of an asr system for dysarthric speaker," *Expert Systems with Applications*, vol. 42, no. 8, pp. 3924–3932, 2015.
- [15] H. Christensen, S. Cunningham, C. Fox, P. Green, and T. Hain, "A comparative study of adaptive, automatic recognition of disordered speech," in *Thirteenth Annual Conference of the International Speech Communication Association*, 2012.
- [16] L. C. Jordan and A. E. Hillis, "Disorders of speech and language: aphasia, apraxia and dysarthria," *Current opinion in neurology*, vol. 19, no. 6, pp. 580–585, 2006.
- [17] A. Kertesz and E. Poole, "The aphasia quotient: the taxonomic approach to measurement of aphasic disability," *Canadian Journal of Neurological Sciences*, vol. 1, no. 1, pp. 7–16, 1974.
- [18] A. Jain, M. Upreti, and P. Jyothi, "Improved accented speech recognition using accent embeddings and multi-task learning," in *Interspeech*, pp. 2454–2458, 2018.
- [19] X. Yang, K. Audhkhasi, A. Rosenberg, S. Thomas, B. Ramabhadran, and M. Hasegawa-Johnson, "Joint modeling of accents and acoustics for multi-accent speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, IEEE, 2018.
- [20] A. Jain, V. P. Singh, and S. P. Rath, "A multi-accent acoustic model using mixture of experts for speech recognition," *Proc. Interspeech 2019*, pp. 779–783, 2019.
- [21] G. Saon, H. Soltau, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," in *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, pp. 55–59, IEEE, 2013.
- [22] B. MacWhinney, D. Fromm, M. Forbes, and A. Holland, "Aphasiabank: Methods for studying discourse," *Aphasiology*, vol. 25, no. 11, pp. 1286–1307, 2011.
- [23] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5329–5333, IEEE, 2018.
- [24] S. Chandrakala, S. Malini, and S. JAYALAKSHMI, "Bag of models based embeddings for assessment of neurological disorders using speech intelligibility," *IEEE Transactions on Emerging Topics in Computing*, 2020.
- [25] J. Williams and S. King, "Disentangling style factors from speaker representations," in *Proc. Interspeech*, vol. 2019, pp. 3945–3949, 2019.
- [26] D. Raj, D. Snyder, D. Povey, and S. Khudanpur, "Probing the information encoded in x-vectors," *arXiv preprint arXiv:1909.06351*, 2019.
- [27] M. Bisani and H. Ney, "Bootstrap estimates for confidence intervals in asr performance evaluation," in *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 1, pp. 1–409, IEEE, 2004.
- [28] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, et al., "The kaldi speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. CONF, IEEE Signal Processing Society, 2011.
- [29] M. Ravanelli, T. Parcollet, and Y. Bengio, "The pytorch-kaldi speech recognition toolkit," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6465–6469, IEEE, 2019.
- [30] R. Gopinath, "Constrained maximum likelihood modeling with gaussian distributions," in *Proceedings Broadcast news transcription and understanding workshop*, 1998.
- [31] D. Povey and G. Saon, "Feature and model space speaker adaptation with full covariance gaussians," in *Ninth International Conference on Spoken Language Processing*, 2006.