

LEVERAGING NATIVE LANGUAGE SPEECH FOR ACCENT IDENTIFICATION USING DEEP SIAMESE NETWORKS

Aditya Siddhant[†], Preethi Jyothi[§], Sriram Ganapathy^{‡*}

[†]Carnegie Mellon University, Pittsburgh, PA, USA

[§]Indian Institute of Technology Bombay, Mumbai, India

[‡]Indian Institute of Science, Bengaluru, India

ABSTRACT

The problem of automatic accent identification is important for several applications like speaker profiling and recognition as well as for improving speech recognition systems. The accented nature of speech can be primarily attributed to the influence of the speaker’s native language on the given speech recording. In this paper, we propose a novel accent identification system whose training exploits speech in native languages along with the accented speech. Specifically, we develop a deep Siamese network based model which learns the association between accented speech recordings and the native language speech recordings. The Siamese networks are trained with i-vector features extracted from the speech recordings using either an unsupervised Gaussian mixture model (GMM) or a supervised deep neural network (DNN) model. We perform several accent identification experiments using the CSLU Foreign Accented English (FAE) corpus. In these experiments, our proposed approach using deep Siamese networks yield significant relative performance improvements of 15.4% on a 10-class accent identification task, over a baseline DNN-based classification system that uses GMM i-vectors. Furthermore, we present a detailed error analysis of the proposed accent identification system.

Index Terms— Accent identification, i-vectors, Deep Siamese networks, Multi-lingual modeling.

1. INTRODUCTION

Over the recent years, many of voice-driven technologies have achieved significant robustness needed for mass deployment. This is largely due to significant advances in automatic speech recognition (ASR) technologies and deep learning algorithms. However, the variability in speech accents pose a significant challenge to state-of-the-art speech systems. In particular, large sections of the English-speaking population in the world face difficulties interacting with voice-driven agents in English due to the mis-match in speech accents

seen in the training data. The accented nature of speech can be primarily attributed to the influence of the speaker’s native language. In this work we focus on the problem of *accent identification*, where the user’s native language is automatically determined from their non-native speech. This can be viewed as a first step towards building accent-aware voice-driven systems.

Accent identification from non-native speech bears resemblance to the task of language identification [1]. However, accent identification is a harder task as many cues about the speaker’s native language are lost or suppressed in the non-native speech. Nevertheless, one may expect that the speaker’s native language is reflected in the acoustics of the individual phones used in non-native language speech, along with pronunciations of words and grammar. In this work, we focus on the acoustic characteristics of an accent induced by a speaker’s native language.

Our main contributions:

- We develop a novel deep Siamese network based model which learns the association between accented speech and native language speech.
- We explore i-vector features extracted using both an unsupervised Gaussian mixture model (GMM) and a supervised deep neural network (DNN) model.
- We present a detailed error analysis of the proposed system which reveals that the confusions among accent predictions are contained within the language family of the corresponding native language.

Section 3 outlines the i-vector feature extraction process. Section 4 describes our Siamese network-based model for accent identification. Our experimental results are detailed in Section 5 and Section 6 provides an error analysis of our proposed approach.

2. RELATED WORK

The prior work on foreign accent identification has drawn inspiration from techniques used in language identification [2].

*This work was carried out with the help of a research grant awarded by Microsoft Research India (MSRI) for the Summer Workshop on Artificial Social Intelligence.

The phonotactic model based approaches [3] and acoustic model based approaches [4] have been explored for accent identification in the past. More recently, i-vector based representations, which is part of the state-of-the-art speaker recognition [5] and language recognition [6] systems, have been applied to the task of accent recognition. The i-vector systems that used GMM-based background models were found to outperform other competitive baseline systems [7, 8, 9].

In the recent years, the language recognition systems and speaker recognition systems have shown promising results with the use of deep neural network (DNN) model based i-vector extraction [10, 11]. Also, none of the previous approaches have exploited speech in native languages while training accent identification systems to the best of our knowledge. This work attempts to develop accent recognition systems using both these components.

3. FACTOR ANALYSIS FRAMEWORK FOR I-VECTOR EXTRACTION

The techniques outlined here are derived from previous work on joint factor analysis (JFA) and i-vectors [12, 13, 14]. We follow the notations used in [12]. The training data from all the speakers is used to train a GMM with model parameters $\lambda = \{\pi_c, \boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c\}$ where π_c , $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$ denote the mixture component weights, mean vectors and covariance matrices respectively for $c = 1, \dots, C$ mixture components. Here, $\boldsymbol{\mu}_c$ is a vector of dimension F and $\boldsymbol{\Sigma}_c$ is assumed to be a diagonal matrix of dimension $F \times F$.

3.1. GMM-based i-vector Representations

Let \mathcal{M}_0 denote the universal background model (UBM) supervector which is the concatenation of $\boldsymbol{\mu}_c$ for $c = 1, \dots, C$ and is of dimensionality $D \times 1$ (where $D = C \cdot F$). Let $\boldsymbol{\Sigma}$ denote the block diagonal matrix of size $D \times D$ whose diagonal blocks are $\boldsymbol{\Sigma}_c$. Let $\mathcal{X}(s) = \{\mathbf{x}_i^s, i = 1, \dots, H(s)\}$ denote the low-level feature sequence for input recording s where i denotes the frame index. Here $H(s)$ denotes the number of frames in the recording. Each \mathbf{x}_i^s is of dimension $F \times 1$.

Let $\mathcal{M}(s)$ denote the recording supervector which is the concatenation of speaker adapted GMM means $\boldsymbol{\mu}_c(s)$ for $c = 1, \dots, C$ for the speaker s . Then, the i-vector model is,

$$\mathcal{M}(s) = \mathcal{M}_0 + \mathbf{V}\mathbf{y}(s) \quad (1)$$

where \mathbf{V} denotes the total variability matrix of dimension $D \times M$ and $\mathbf{y}(s)$ denotes the i-vector of dimension M . The i-vector is assumed to be distributed as $\mathcal{N}(\mathbf{0}, \mathbf{I})$.

In order to estimate the i-vectors, the iterative EM algorithm is used. We begin with a random initialization of the total variability matrix \mathbf{V} . Let $p_\lambda(c|\mathbf{x}_i^s)$ denote the alignment probability of assigning the feature vector \mathbf{x}_i^s to mixture com-

ponent c . The sufficient statistics are then computed as,

$$\begin{aligned} N_c(s) &= \sum_{i=1}^{H(s)} p_\lambda(c|\mathbf{x}_i^s) \\ \mathbf{S}_{X,c}(s) &= \sum_{i=1}^{H(s)} p_\lambda(c|\mathbf{x}_i^s)(\mathbf{x}_i^s - \boldsymbol{\mu}_c) \end{aligned} \quad (2)$$

Let $\mathbf{N}(s)$ denote the $D \times D$ block diagonal matrix with diagonal blocks $N_1(s)\mathbf{I}, N_2(s)\mathbf{I}, \dots, N_C(s)\mathbf{I}$ where \mathbf{I} is the $F \times F$ identity matrix. Let $\mathbf{S}_X(s)$ denote the $D \times 1$ vector obtained by splicing $\mathbf{S}_{X,1}(s), \dots, \mathbf{S}_{X,C}(s)$.

It can be easily shown [12] that the posterior distribution of the i-vector $p_\lambda(\mathbf{y}(s)|\mathcal{X}(s))$ is Gaussian with covariance $\mathbf{l}^{-1}(s)$ and mean $\mathbf{l}^{-1}(s)\mathbf{V}^*\boldsymbol{\Sigma}^{-1}\mathbf{S}_X(s)$, where

$$\mathbf{l}(s) = \mathbf{I} + \mathbf{V}^*\boldsymbol{\Sigma}^{-1}\mathbf{N}(s)\mathbf{V} \quad (3)$$

The optimal estimate for the i-vector $\mathbf{y}(s)$ obtained as $\text{argmax}_{\mathbf{y}} [p_\lambda(\mathbf{y}(s)|\mathcal{X}(s))]$ is given by the mean of the posterior distribution.

For re-estimating the \mathbf{V} matrix, the maximization of the expected value of the log-likelihood function (EM algorithm), gives the following relation [12],

$$\sum_{s=1}^S \mathbf{N}(s) \mathbf{V} \mathbb{E}[\mathbf{y}(s)\mathbf{y}^*(s)] = \sum_{s=1}^S \mathbf{S}_X(s) \mathbb{E}[\mathbf{y}^*(s)] \quad (4)$$

where $\mathbb{E}[\cdot]$ denotes the posterior expectation operator. The solution for Eq. (4) can be computed for each row of \mathbf{V} . Thus, the i-vector estimation is performed by iterating between the estimation of posterior distribution and the update of the total variability matrix (Eq. (4)).

3.2. DNN i-vectors

Instead of using a GMM-UBM based computation of i-vectors, we can also use DNN-based context dependent state (senone) posteriors to generate the sufficient statistics used in the i-vector computation [15, 10]. The GMM mixture components will be replaced with the senone classes present at the output of the DNN. Specifically, $p_\lambda(c|\mathbf{x}_i^s)$ used in Eq. (2) is replaced with the DNN posterior probability estimate of the senone c given the input acoustic feature vector \mathbf{x}_i^s and the total number of senones is the parameter C . The other parameters of the UBM model $\lambda = \{\pi_c, \boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c\}$ are computed as

$$\begin{aligned} \pi_c &= \frac{\sum_{s=1}^S \sum_{i=1}^{H(s)} p(c|\mathbf{x}_i^s)}{\sum_{c=1}^C \sum_{s=1}^S \sum_{i=1}^{H(s)} p(c|\mathbf{x}_i^s)} \\ \boldsymbol{\mu}_c &= \frac{\sum_{s=1}^S \sum_{i=1}^{H(s)} p(c|\mathbf{x}_i^s)\mathbf{x}_i^s}{\sum_{c=1}^C \sum_{s=1}^S \sum_{i=1}^{H(s)} p(c|\mathbf{x}_i^s)} \\ \boldsymbol{\Sigma}_c &= \frac{\sum_{s=1}^S \sum_{i=1}^{H(s)} p(c|\mathbf{x}_i^s)(\mathbf{x}_i^s - \boldsymbol{\mu}_c)(\mathbf{x}_i^s - \boldsymbol{\mu}_c)^*}{\sum_{c=1}^C \sum_{s=1}^S \sum_{i=1}^{H(s)} p(c|\mathbf{x}_i^s)} \end{aligned} \quad (5)$$

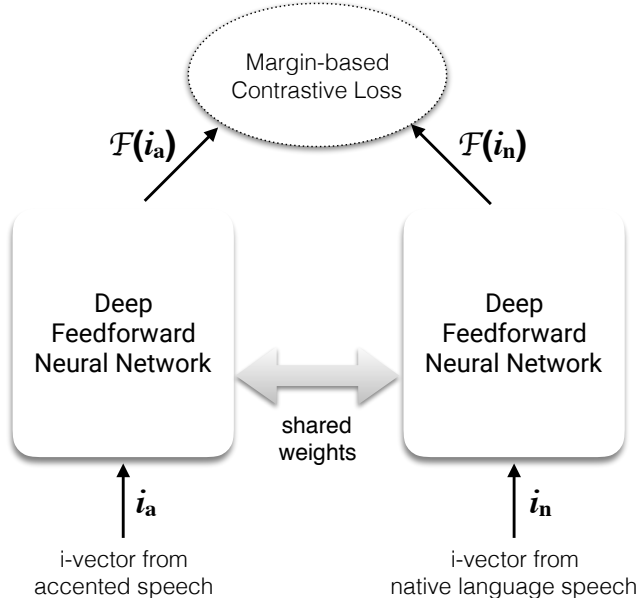


Fig. 1. Siamese network architecture for accent identification

4. PROPOSED APPROACH

Siamese networks [16, 17] are neural network models that are designed to learn a similarity function between pairs of input representations. This architecture consists of two identical neural networks with shared weights where the input is a pair of samples. The objective is to find a function that minimizes or maximizes the similarity between the pair of inputs depending on whether they belong to the same category or not. This is achieved by optimizing a contrastive loss function containing dual terms - a component that reduces the contribution from the positive training examples (i.e. pairs of inputs belonging to the same category) and a component that increases the contribution from negative training examples (i.e. pairs of inputs from different categories).

Figure 1 shows an illustration of the Siamese network we used for accent identification. Each training example comprises a pair of input i-vectors, $\{\mathbf{i}_1, \mathbf{i}_2\}$ corresponding to an accented speech sample and a language speech sample, and a binary label $y \in \{0, 1\}$ indicating whether or not the native language underlying the accented speech sample exactly matches that of the language speech sample. The positive training examples correspond to accented speech i-vectors that are paired with native language i-vectors from the same speaker. For the negative examples, we paired up accented speech i-vectors with i-vectors from languages different from the one underlying the accented speech sample. These training instances are fed to twin networks with shared parameters which produce two outputs $\{\mathcal{F}(\mathbf{i}_1), \mathcal{F}(\mathbf{i}_2)\}$ corresponding to the input i-vectors \mathbf{i}_1 and \mathbf{i}_2 . The whole network is then

LANGUAGE	Accented speech			Native language speech
	Training	Dev	Test	
BP	92	30	31	198
HI	66	22	22	206
FA	50	17	16	182
GE	55	18	18	161
HU	51	17	17	187
IT	34	11	12	168
MA	52	18	18	189
RU	44	14	14	172
SP	31	9	10	140
TA	37	13	12	128

Table 1. Statistics of accented English and native language speech data. All the displayed numbers correspond to minutes of speech.

trained to minimize the following contrastive loss function:

$$\mathcal{L}(\mathbf{i}_1, \mathbf{i}_2, y) = (1 - y) \cdot (d(\mathcal{F}(\mathbf{i}_1), \mathcal{F}(\mathbf{i}_2)))^2 + y \cdot (\max(0, 1 - d(\mathcal{F}(\mathbf{i}_1), \mathcal{F}(\mathbf{i}_2))))^2 \quad (6)$$

where $d(\cdot, \cdot)$ is a distance function between the output representations.

We use a large number of positive and negative training samples to learn a distance metric between the accented speech i-vectors and the language i-vectors. During test time, we compare the accented speech test i-vector with a representative language i-vector and choose the language whose i-vector representations are the nearest from the accented speech i-vector, according to the distance metric learned by the Siamese network. We experiment with different strategies to determine how the language i-vectors should be constructed during test time. Section 5.3 discusses more details of these test strategies.

5. EXPERIMENTAL RESULTS

5.1. Data

For our experiments, we used the CSLU Foreign Accented English Release 1.2 database [18] that consists of telephone-quality spontaneous English speech by native speakers of 22 different languages. We set up a 10-class accent identification task using accented English from speakers of 10 different languages which had the most amount of data: Brazilian Portuguese (BP), Hindi (HI), Farsi (FA), German (GE), Hungarian (HU), Italian (IT), Mandarin Chinese (MA), Russian (RU), Spanish (SP) and Tamil (TA). For native language speech, we used the CSLU 22 Languages Corpus [19] which contains telephone-quality continuous speech in all the above-mentioned 10 languages. Many of the speakers in the CSLU 22 Languages corpus also recorded speech samples for the CSLU Foreign Accented English corpus. The samples from

Classifier	GMM i-vectors		DNN i-vectors	
	Dev	Test	Dev	Test
LDA	33.5	37.2	39.8	43.8
SVM	35.1	40.2	39.8	45.2
NNET	35.8	40.8	41.4	44.8

Table 2. Accuracy rates (%) from classifiers using both GMM i-vectors and DNN i-vectors.

these speakers were used to construct positive examples for training our Siamese network. Table 1 gives detailed statistics about the data used in our experiments, along with the training, development and test set splits. These splits were created using a stratified random sampler so that the proportion of different accents in each set remains almost same.

For the GMM based i-vectors, a 2048-component UBM was trained and 400 dimensional i-vectors were extracted using the formulation given in Sec. 3. These features will be referred to as *GMM i-vectors* in the rest of the paper. The UBM was trained with 39 dimensional MFCC features which were mean and variance normalized at the utterance level. The training data used in the UBM was obtained from the multilingual corpora from NIST SRE 2008 (consisting of telephone-quality speech) and the Switchboard English database [20]. For training the DNN i-vectors, an acoustic model was developed for Switchboard using the Kaldi toolkit [21]. The DNN model generates 4677-dimensional senone posterior features which are used in the i-vector extraction (Sec. 3.2). The i-vector training based on the DNN-UBM uses data from NIST SRE08 and Switchboard. We use 300 dimensional i-vectors from the DNN model (henceforth referred to as *DNN i-vectors*). Both i-vectors were length normalized before the classifier training.

Performance evaluation: We used accent identification accuracy as the primary metric to evaluate our proposed approach. This is computed as the percentage of utterances which are correctly identified as having one of the 10 above-mentioned accents. (A classifier based on chance would give an accuracy of 10% on this task.)

5.2. Comparing GMM i-vectors with DNN i-vectors

Table 2 shows the performance of various baseline accent identification systems using both GMM i-vectors and DNN i-vectors as input features. We used an LDA-based classifier which reduces the dimensionality of the input vectors by linear projection onto a lower dimensional space that maximizes the separation between classes. We also built an SVM classifier using a radial basis function (RBF) kernel. Finally, NNET is a feed-forward neural network (with a single 128-node hidden layer) that is trained to optimize categorical cross-entropy using the Adam optimization algorithm [22]. LDA and SVM were implemented using the scikit-learn library [23] and

System	Dev Accuracy	Test Accuracy
Siamese-1	42.7	46.4
Siamese-2	43.3	46.8
Siamese-3	43.3	47.3
Siamese-4	43.6	47.9

Table 3. Performance (%) of Siamese network-based classifier using different test strategies.

NNET was implemented using the Keras toolkit [24]. The hyper-parameters of all these models were tuned using the validation set. From Table 2, we observe that the classifiers using DNN i-vectors clearly outperform the classifiers using the GMM i-vectors. This is intuitive because the DNN i-vectors carry more information about the underlying phones. Both the SVM and NNET classifiers perform comparably and outperformed the linear (LDA) classifier due to the inherent non-linearity in the data.

In all subsequent experiments, we use the 300-dimensional DNN i-vectors (unless mentioned otherwise).

5.3. Evaluating the Siamese network

We tried various configurations of the Siamese networks, along with varying ratios of positive and negative training examples. After tuning on the validation set, the Siamese architecture which yielded the best result consisted of 2 hidden layers with 128 nodes each and a dropout rate of 0.2 in the first hidden layer [25]. We used the RMSprop optimizer and the Glorot initialization for the network [26]. The network was trained on 100,000 positive and 900,000 negative training pairs. The 10 accents were equally distributed across the positive and negative pairs.

Table 3 lists the accuracies of the Siamese network-based classifiers using different strategies to choose language i-vectors during test time. We first extracted a random sample of 30 language i-vectors and computed the mean of the lowest five output scores. (Here, 30 and 5 were tuned on the validation set.) The language that was least distant from the accented test sample was chosen as the predicted output. This is the system referred to as “Siamese-1”. “Siamese-2” refers to a system where we computed a mean language i-vector across all the i-vectors for a particular language. For “Siamese-3”, we first clustered the language i-vectors into 4 clusters using the k-means algorithm, following which we computed cluster means and chose the language whose cluster mean was minimally distant from the accented test sample. Finally, “Siamese-4” augments “Siamese-3” with a neural network classifier (consisting of two 8-node hidden layers with 40 input nodes and 10 output nodes). This second DNN is trained on the distance measures computed from Siamese-3 model for the 10 accent classes (4 scores for each file obtained from the four cluster mean vectors) and it predicts the target accent class. This network is trained on the

Classifier	GMM i-vectors		DNN i-vectors	
	Dev	Test	Dev	Test
SVM	35.1	40.2	39.8	45.2
NNET	35.8	40.8	41.4	44.8
Siamese-4	37.8	42.3	43.6	47.9

Table 4. Performance (accuracy %) of Siamese network-based classifier compared to baseline systems.

System	Dev Accuracy	Test Accuracy
NNET-append	38.6	41.1
NNET-nonid-twin	41.9	44.8
NNET-transfer	41.7	45.3
Siamese-4	43.6	47.9

Table 5. Performance of various classifiers that use native language speech during training.

validation data. In our experiments, “Siamese-4” provided the best performance.

Table 4 compares the performance of our best-performing Siamese network to the two best-performing baseline systems from Table 2. We see consistent improvements from using the Siamese network classifier over the best baseline system on both the validation set and the test set. These improvements hold when both GMM i-vectors and DNN i-vectors are used as input features.

5.4. Comparison with other systems using native language i-vectors

All the baseline systems used so far only made use of the accented English samples during training and did not make use of the native language speech samples. We compare our Siamese network-based approach with other techniques that exploit speech data from native languages during training. First, analogously to “NNET”, we train a 2-layer feed-forward neural network but with input features consisting of language i-vectors concatenated with accent i-vectors. This system is referred to as “NNET-append”. We build a second system, which we call “NNET-nonid-twin”, that is identical to the twin network Siamese architecture shown in Figure 1, except the weights of the twin networks are not shared. Finally, we also investigate a transfer learning based approach, referred to as “NNET-transfer”. For this system, we train a 2-layer feed-forward neural network using only language i-vectors to predict the underlying language. Then, we use the resulting weights from the hidden layers as an initialization for a neural network that uses accent i-vectors as inputs to predict the underlying accent. Table 5 compares these three systems with “Siamese-4” introduced in Table 3. We observe that our proposed Siamese-network system outperforms all the other systems which also have access to native language i-vectors.

Accent judgment (1-4)	% of samples	Accuracy
1	10	34.7
2	10	41.3
3	79	50.4
4	1	56.2

Table 6. Accuracies on utterances of varying accent strengths.

5.5. Accuracies on accented speech utterances with varying accent strengths

The CSLU Foreign Accented Speech corpus contains perceptual judgments about the accents in the utterances. Each accented speech sample was independently annotated for accent strength on a scale of 1-4 (where 1 denotes a very mild accent and 4 denotes a very strong accent) judged by three native American English speakers. Table 6 shows how the Siamese network-based classifier (“Siamese-4”) performs on utterances when grouped according to accent strength. Intuitively, as the accent strength increases, our classifier accuracy increases. Despite the fact that our test set predominantly contains utterances of accent strength 3, it is interesting to see that the average accuracy on these test samples is much higher than the samples rated 1 and 2 on accent strength.

5.6. System Combination

Table 7 shows the accuracies from our two top baseline systems and our best Siamese network based system, Siamese-4, when the correct accent is among the top 2 and top 3 accent predictions. We observe that the accuracies dramatically improve across all three individual systems when moving from 1-best to 2-best and 3-best accuracies. This indicates that a significant part of the confusions seen across the classes are confined to the top 3 predictions.

We also combine the outputs from the three individual systems. We adopt a simple majority voting strategy: Choose the prediction that 2 or more systems agree upon. If all three systems predict different accent classes, then choose the accent predicted by Siamese-4. We also use a weighted voting

System	1-best	2-best	3-best
Individual systems			
SVM	45.2	64.9	73.1
NNET	44.8	64.1	72.4
Siamese-4	47.9	70.2	80.4
Combined systems			
Majority-voting	48.6	73.3	85.1
Weighted-voting	48.8	75.1	87.4

Table 7. N-best accuracies (%) on the test set from both individual systems and combined systems.

strategy. The outputs from Siamese-4 were converted into posterior probabilities using a softmax layer. A weighted combination of the probabilities from SVM, NNET and Siamese-4 was used to determine the accent with the highest probability. (Weights of 0.3, 0.3 and 0.4 were used for SVM, NNET and Siamese-4, respectively.) The N-best accuracies for both these combined systems are shown in Table 7. As expected, we observe performance gains from combining the outputs of all three individual systems; the gains are larger in the 2-best and 3-best cases.

6. DISCUSSION

It is illustrative to visualize the accent predictions made by our proposed Siamese network-based classifier in order to learn more about the main confusions that are prevalent in the system. We visualize the confusion matrix of the Siamese-4 classifier on the test samples using a heat map as shown in Figure 2. Darker/bigger bubbles against a column for a given row indicate a larger number of samples were predicted as being of the accent labeled against the column. For each accent class, Figure 2 shows there are a sizable number of test examples that are correctly predicted as evidenced by the dark bubbles along the diagonal. We also see more than one darkly-colored bubble along the non-diagonal indicating strong evidence for confusion. For example, Hindi-accented English samples are often confused as being Tamil-accented and conversely, Tamil-accented in some test samples are mistaken for Hindi-accented. This is very intuitive given that both these accents are very closely related and correspond to the same geographical region. Indeed, if we group the languages according to the language families that they belong to, i.e. {BP, RU and IT}, {SP, GE and HU}, {MA}, {HI and TA} and {FA}, the corresponding confusion matrix exhibits far less confusion as shown in Figure 3.

7. CONCLUSIONS

In this paper, we explore the problem of accent identification from non-native English speech. We propose a novel approach based on deep Siamese neural networks that uses i-vectors extracted from both accented speech and native language speech samples and learns a semantic distance between these feature representations. On a 10-class accent identification task, our proposed approach outperforms a neural network-based classifier using both GMM-based i-vectors and DNN-based i-vectors with relative accuracy improvements of 15.4% and 7.0%, respectively.

In this work, we focused on the acoustic characteristics of an accent induced by a speaker’s native language. Accents are also correlated with specific lexical realizations of words in terms of pronunciations and variations in word usage and grammar. As future work, we plan to explore how to incorpo-




Fig. 2. Bubble plot visualizing the confusion matrix of test set predictions from Siamese-4.




Fig. 3. Bubble plot visualizing the confusion matrix of test set predictions from Siamese-4, after grouping related languages.

rate the pronunciation model and language model based features to automatic identification of speech accents.

8. ACKNOWLEDGEMENTS

The authors thank the organizers of the MSR-ASI workshop, Monojit Choudhury, Kalika Bali and Sunayana Sitaram for the opportunity and all their help, as well as the other project team members Abhinav Jain, Gayatri Bhat and Sakshi Kalra for their support. We also gratefully acknowledge the logistical support from Microsoft Research India (MSRI) for this project as well as access to Microsoft Azure cloud computing resources.

9. REFERENCES

- [1] Marc A Zissman and Kay M Berkling, "Automatic language identification," *Speech Communication*, vol. 35, no. 1, pp. 115–124, 2001.
- [2] Marc A Zissman, "Comparison of four approaches to automatic language identification of telephone speech," *IEEE Transactions on speech and audio processing*, vol. 4, no. 1, pp. 31, 1996.
- [3] F. Biadsy, J. Hirschberg, and D. P. W. Ellis, "Dialect and accent recognition using phonetic-segmentation super-vectors," in *Proceedings of Interspeech*, 2011.
- [4] Carlos Teixeira, Isabel Trancoso, and António Serralheiro, "Accent identification," in *Proceedings of ICSLP*. IEEE, 1996, pp. 1784–1787.
- [5] Najim Dehak, Patrick J Kenny, Réda Dehak, Pierre Dumouchel, and Pierre Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.
- [6] Najim Dehak, Pedro A Torres-Carrasquillo, Douglas Reynolds, and Reda Dehak, "Language recognition via i-vectors and dimensionality reduction," in *Proceedings of Interspeech*, 2011.
- [7] Mohamad Hasan Bahari, Rahim Saeidi, Hugo Van hamme, and David Van Leeuwen, "Accent recognition using i-vector, gaussian mean supervector and gaussian posterior probability supervector for spontaneous telephone speech," in *Proceedings of ICASSP*. IEEE, 2013, pp. 7344–7348.
- [8] Alexandros Lazaridis, Elie Khoury, Jean-Philippe Goldman, Mathieu Avanzi, Sébastien Marcel, and Philip N Garner, "Swiss french regional accent identification," in *Proceedings of Odyssey*, 2014.
- [9] Maryam Najafian, Saeid Safavi, Phil Weber, and Martin Russell, "Identification of British English regional accents using fusion of i-vector and multi-accent phonotactic systems," in *Proceedings of Odyssey*, 2016.
- [10] Yun Lei, Luciana Ferrer, Aaron Lawson, Mitchell McLaren, and Nicolas Scheffer, "Application of convolutional neural networks to language identification in noisy conditions," in *Proc. Speaker Odyssey Workshop*, 2014.
- [11] Seyed Omid Sadjadi, Sriram Ganapathy, and Jason W Pelecanos, "The IBM 2016 speaker recognition system," *Proceedings of Odyssey*, 2016.
- [12] Patrick Kenny, Gilles Boulianne, and Pierre Dumouchel, "Eigenvoice modeling with sparse training data," *IEEE Transactions on Speech and Audio Processing*, vol. 13, no. 3, pp. 345–354, 2005.
- [13] Patrick Kenny, "Joint factor analysis of speaker and session variability: Theory and algorithms," *CRIM, Montreal, (Report) CRIM-06/08-13*, 2005.
- [14] Najim Dehak, Patrick Kenny, Réda Dehak, Pierre Dumouchel, and Pierre Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.
- [15] Mitchell McLaren, Yun Lei, Nicolas Scheffer, and Luciana Ferrer, "Application of convolutional neural networks to speaker recognition in noisy conditions," in *Proceeding of Interspeech*, 2014.
- [16] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah, "Signature verification using a Siamese time delay neural network," in *Proceedings of NIPS*, 1994, pp. 737–744.
- [17] Sumit Chopra, Raia Hadsell, and Yann LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *Proceedings of CVPR*. IEEE, 2005, vol. 1, pp. 539–546.
- [18] T. Lander, "CSLU: Foreign Accented English Corpus Release 1.2," <https://catalog.ldc.upenn.edu/ldc2007s08>, 2007.
- [19] T. Lander, "CSLU: 22 Languages Corpus," <https://catalog.ldc.upenn.edu/ldc2005S26>, 2005.
- [20] John J Godfrey, Edward C Holliman, and Jane McDaniel, "SWITCHBOARD: telephone speech corpus for research and development," in *Proceedings of ICASSP*, 1992, pp. 517–520.
- [21] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely, "The Kaldi speech recognition toolkit," in *Proceedings of ASRU*, 2011.
- [22] Diederik Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [23] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos,

D. Courvapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

- [24] François Chollet et al., “Keras,” <https://github.com/fchollet/keras>, 2015.
- [25] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [26] Xavier Glorot and Yoshua Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *Proceedings of AISTATS*, 2010, pp. 249–256.