HKUST SPD - INSTITUTIONAL REPOSITORY

Title

Authors

Source

Version

DOI

Publisher

Copyright

This version is available at HKUST SPD - Institutional Repository (<https://repository.ust.hk/ir>)

If it is the author's pre-published version, changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published version.

Eigentrigraphemes for Under-Resourced Languages

Tom Ko and Brian Mak

*Department of Computer Science and Engineering,
The Hong Kong University of Science and Technology,
Clear Water Bay, Hong Kong.
E-mail: {tomko, mak}@cse.ust.hk.*

Abstract

Grapheme-based modeling has an advantage over phone-based modeling in automatic speech recognition for under-resourced languages when a good dictionary is not available. Recently we proposed a new method for parameter estimation of context-dependent hidden Markov model (HMM) called *eigentriphone modeling*. Eigentriphone modeling outperforms conventional tied-state HMM by eliminating the quantization errors among the tied states. The eigentriphone modeling framework is very flexible and can be applied to any group of modeling unit provided that they may be represented by vectors of the same dimension. In this paper, we would like to port the eigentriphone modeling method from a phone-based system to a grapheme-based system; the new method will be called *eigentrigrapheme modeling*. Experiments on four official South African under-resourced languages (Afrikaans, South African English, Sesotho, siSwati) show that the new eigentrigrapheme modeling method reduces the word error rates of conventional tied-state trigrapheme modeling by an average of 4.08% relative.

Keywords: Eigentriphone, eigentrigrapheme, under-resourced language, grapheme, regularization, weighted PCA.

---

1. Introduction

In the past, research efforts on automatic speech recognition (ASR) are highly focused on the most popular languages such as English, Mandarin, Japanese, French, German, . . . etc. in the developed countries. The remaining languages in the world, lacking audio and language resources, are considered under-resourced languages. Usually the phonetics and linguistics of these languages are not well studied either, thus the development of human language technologies for these languages is greatly hindered. Nonetheless, some of these under-resourced languages are spoken by a large population. For example, Vietnamese is spoken by about 80 million people, and Thai is spoken by 60 million people. It is not difficult to see that real-life ASR applications for these languages have a great potential. One major obstacle in developing an ASR system for under-resourced

languages is the availability of data. It is usually costly and labor-intensive to create audio recordings and their human annotated transcriptions, and make linguistic analyses for languages. As a consequence, it is both academically intriguing and commercially attractive to look for economically more efficient and faster ways to create human language technologies for the under-resourced languages.

In order to reduce the amount of annotated audio data for training the acoustic models of a new target language, cross-lingual [1, 2] and multi-lingual [3] acoustic modeling techniques have been developed. The rationale behind these techniques is that an acoustic model may be ported to or adapted from some other high-resourced languages, and only a relatively small amount of training data is required for the target language. A key step for these cross-lingual or multi-lingual techniques to work is to figure out a good mapping between phonemes across the languages. This can be done either using a knowledge-based [4] or data-driven approach [3]. In the data-driven approach, the similarities between sounds can be measured by various distance measures such as confusion matrix [4], entropy-based distance [3] or Euclidean distance [5]. The approach is further improved when the underlying model is more compactly represented. A notable example is the use of subspace Gaussian mixture model (SGMM) [6] in multi-lingual ASR [7, 8]. Another research direction is heading to making linguistic analysis of a target language easier and faster. Deducing the phone set and preparing the pronunciation dictionary for a new language usually require native linguistic experts. This process is expensive and time-consuming, and is even more so for non-native developers. One way to partially automate the development of a pronunciation dictionary is to first prepare a small primary dictionary manually, and then use it to bootstrap a large dictionary by applying grapheme-to-phoneme conversion [9, 10, 11, 12]. However, the performance of the final dictionary highly depends on the quality of the primary one. If the primary dictionary is not rich enough and does not cover all the implicit grapheme-to-phoneme relations in the language, the performance of the overall system will be hampered.

On the other hand, there is a simple solution to the creation of the phone set and pronunciation dictionary for an under-resourced language: there is no need to develop them if graphemes instead of phonemes are adopted as the acoustic modeling units. In grapheme modeling [13, 14, 15, 2], each word in the “pronunciation dictionary” is simply represented by its graphemic transcription according to its lexical form. According to [16], there are six types of writing systems in the world: logosyllabary, syllabary, abjad, alphabet, abugida, and featural. Many languages that use the alphabet writing system are suitable for grapheme acoustic modeling, and their grapheme set is usually selected to be the same as their alphabet set [13].

The performance of grapheme modeling in ASR is sensitive to the languages. For example, it works better than phone modeling in Spanish but worse than phone modeling in English and Thai [17]. The reason is that the pronunciation of English has developed away from its written form over time, whereas Thai has some complex rules that map its writing to the pronunciation. There are

techniques that improve grapheme modeling; for example, in [15], a text normalization scheme was applied on Thai graphemes to improve the performance of a Thai ASR system. There are also works on multi-lingual grapheme modeling [18, 19]. These techniques, however, are usually language-dependent as linguistic knowledge of the target language has to be known in advance. In this paper, we will investigate a language-independent technique to improve current grapheme modeling.

Recently, we proposed a new method of estimating parameters of context-dependent models called *eigentriphone* acoustic modeling [20, 21, 22, 23]. The main idea behind our method is the derivation of an eigenbasis over a set of triphones so that each triphone in the set can then be modeled as a distinct point in the space spanned by the basis vectors. The basis vectors are now called *eigentriphones* in our method. The eigentriphones represent the most important context-dependent characteristics among the given set of triphones. Since usually not many eigentriphones are required to represent the eigenspace, triphones with few training samples can be robustly estimated as a linear combination of the eigentriphones. From another point of view, context-dependent phonetic information is extracted from the more frequently occurring triphones in the triphone set in the form of basis vectors, which are shared with the less frequently occurring triphones in the set. The acoustic modeling of those triphones with limited amount of training data may then be thought of as an adaptation problem which is then solved by the eigenvoice approach [24]. Moreover, compared with conventional tied-state triphone modeling, our new eigentriphone modeling method can eliminate the inevitable quantization error due to state tying — states tied together are not distinguishable. In our previous works, we investigated the use of model-based, state-based, and cluster-based eigentriphone acoustic modeling, and observed that cluster-based eigentriphone modeling consistently outperformed the conventional tied-state hidden Markov model (HMM) training method on TIMIT phoneme recognition and WSJ word recognition [23].

Eigentriphone acoustic modeling is complementary to other acoustic modeling approaches such as SGMM that are based on tied-state HMM-GMMs. SGMM allows a compact representation of HMMs with a set of common basis Gaussians and state-dependent projection vectors. Since there are fewer parameters to estimate, model parameters of HMM-SGMM may be estimated more robustly for the same amount of training data. From another perspective, HMM-SGMMs may be trained with less data than HMM-GMMs. Whereas state are usually tied in HMM-SGMMs, eigentriphone modeling unties the tied states so that the states in the final models are, in general, all distinct. Eigentriphone modeling may take HMM-GMMs or HMM-SGMMs as its initial models, unties the states in each state cluster, derives an eigenbasis in each cluster, and re-estimates the mean vectors of each member state in the state cluster as a combination of the eigenvectors. As a result, eigentriphone modeling recovers the quantization errors in the initial tied-state HMM-GMMs or HMM-SGMMs so that the final models are more discriminative.

Although we call our method eigentriphone acoustic modeling, the proposed

framework is actually very flexible and can be applied to any group of modeling units provided that they may be represented by vectors of the same dimension. In this paper, we would like to port the *cluster-based eigentriphone modeling* framework to a grapheme-based ASR system to estimate the parameters of trigrapheme acoustic models. The new method, which we call *cluster-based eigentrigrapheme acoustic modeling*, have the following favorable properties over other grapheme-based methods:

- Since it uses graphemes as the modeling units, it enjoys the same benefits that other grapheme-based modeling methods do. Most importantly, there is no need to create a phone set and a pronunciation dictionary. Thus, it is more favorable for building ASR system for under-resourced languages.
- Eigentrigrapheme modeling will also enjoy the same benefit like eigentriphone modeling: Many trigraphemes in under-resourced languages may have little training data; in the past, the problem is mainly solved by state tying, but eigentrigrapheme modeling allows reliable estimation of the infrequently occurring trigraphemes by careful state clustering and then projecting the member states of each cluster onto a low-dimensional subspace spanned by a small set of eigentrigraphemes of the cluster.
- No language-specific knowledge is required and the whole method is data-driven. It can be used to improve existing systems that are based on conventional tied-state trigrapheme HMM. In fact, one may implement our method as a post-processing procedure on conventional tied-state trigrapheme HMMs.
- If trigrapheme state clusters are created using graphemic decision tree, the decision tree may also be used to synthesize unseen trigraphemes in the test lexicon.
- States are generally not tied among the final trigrapheme models. Thus, the states are generally distinct from each other. We believe that the final trigrapheme models can be more discriminative than those trained on other methods.

The rest of this paper is organized as follows. In Section 2, we will describe the cluster-based eigentrigrapheme acoustic modeling method. That is followed by experimental evaluation in Section 3 and conclusions in Section 4.

2. Cluster-based Eigentrigrapheme Acoustic Modeling

Fig. 1 shows an overview of the cluster-based eigentrigrapheme acoustic modeling method. The framework is very similar to the cluster-based eigentriphone modeling method [23]. All trigrapheme states are first represented by some

Figure 1: The cluster-based eigentrigrapheme acoustic modeling method. (WPCA = weighted principal component analysis; PMLED = penalized maximum-likelihood eigen-decomposition)

supervectors and they are assumed to lie in a low dimensional space¹ spanned by a set of eigenvectors. In other words, each trigrapheme supervector is a linear combination of a small set of eigenvectors which are now called eigentrigraphemes. Clustering of the states can be done by a singleton decision tree, and the procedure is exactly the same as that of creating a conventional tied-state trigrapheme system.

Cluster-based eigentrigrapheme modeling consists of three major steps: (a) state clustering via a singleton decision tree, (b) derivation of the eigenbasis, and (c) estimation of eigentrigrapheme coefficients. They will be discussed in further details below.

2.1. Trigrapheme State Clustering (or Tying) by a Singleton Decision Tree

One major difference between phone-based and grapheme-based acoustic modeling lies in the construction of the decision tree for tying hidden Markov model (HMM) states. In phone-based modeling, it is well-known that decision tree using phonetic questions [25] can significantly improve speech recognition performance by striking a good balance between trainability and resolution of the acoustic models. However, it is not clear how the phonetic questions used in a phone-based system to tie triphone states can be ported to tie trigrapheme

¹The dimension of the space is low when compared with the dimension of the trigrapheme state supervectors.

states in a grapheme-based system as the relation between the graphemes and their influence in the pronunciation of their neighboring graphemes is not well understood. Although the questions may be re-defined manually [14] or automatically, [13] investigated the performance of both methods in several languages and concluded that questions asking only the identity of the immediate neighboring grapheme, named as *singleton questions*, work at least as well as other types of questions. In this paper, decision tree² using singleton questions at each node is used to generate the conventional tied-state trigram HMMs. In addition, the trigram states that belong to the same tied state naturally form a state cluster on which our new cluster-based eigentrigram modeling may be applied. In other words, the same singleton decision tree can be used to create the tied states for a conventional tied-state trigram system as well as the state clusters for the construction of cluster-based eigentrigrams³.

2.2. Conventional Tied-state Trigram HMM Training

We adopt the standard procedure in HTK [26] to create the conventional tied-state trigram HMM system as follows.

STEP 1 : Context-independent grapheme acoustic models are estimated from the training data. Each context-independent grapheme model is a 3-state strictly left-to-right HMM, and each state is represented by a single Gaussian.

STEP 2 : Each context-independent grapheme HMM is then cloned to initialize *all* its context-dependent trigrams.

STEP 3 : For each base grapheme, the transition probabilities of all its trigrams are tied together.

STEP 4 : For each base grapheme, tie the corresponding HMM states of all its trigrams using a singleton decision tree. Thus, three singleton decision trees are built for each base grapheme. Once a set of trigram states are tied together, they share the same set of Gaussian means, diagonal covariances, and mixture weights.

STEP 5 : Synthesize the unseen trigrams by going through the singleton questions of the decision trees.

STEP 6 : Grow the Gaussian mixtures of the models⁴ with the training data until each tied state is represented by an M -component Gaussian mixture

---

²The questions in the decision tree are generated from the grapheme set of the target language which is derived by scanning through the training data. Thus, the trees are language-dependent but our method is still language-independent.

³Although the state clusters of cluster-based eigentrigram modeling and the tied states of conventional trigram modeling both come from the nodes of the same decision tree, in general, they may not be exactly the same nodes. The optimal set of tied states or state clusters is determined using a separate set of development speech data.

⁴In this paper, 4 iterations of embedded Baum-Welch training was applied every time after the number of mixtures was grown.

model (GMM) with diagonal covariance. In practice, the optimal value of M is determined by a separate set of development data.

2.3. Eigentrigrapheme Acoustic Modeling

Recall that each node in the state clustering decision tree has dual roles: it is treated as a tied state for tied-state HMM training, and as a state cluster for eigentrigrapheme modeling. To begin cluster-based eigentrigrapheme modeling, one first decides the tree nodes to be used as the state clusters. Then the state-clusters are treated as tied states, and conventional tied-state trigrapheme HMMs are created using the procedure described in Section 2.2. The resulting tied-state HMMs are used as the initial models for deriving the eigentrigraphemes of each state cluster.

2.3.1. Derivation of Cluster-based Eigentrigraphemes

The following procedure is repeated for each state cluster i , consisting of N_i member states.

STEP 7 : Untie the Gaussian means of all the trigrapheme states in a state cluster with the exception of the unseen trigrapheme states. The means of the cluster GMM are then cloned to initialize *all* untied trigrapheme states in the cluster. Note that the Gaussian covariances and mixture weights of the states in the cluster are still tied together.

STEP 8 : Re-estimate only the Gaussian means of trigrapheme states after cloning. Their Gaussian covariances and mixture weights remain unchanged as those of their state cluster GMM.

STEP 9 : Create a trigrapheme state supervector $\mathbf{v}_{ip}$ for each trigrapheme state p in state cluster i by stacking up all its Gaussian mean vectors from its M -component GMM as below

$$\mathbf{v}_{ip} = [\boldsymbol{\mu}_{ip1}, \boldsymbol{\mu}_{ip2}, \dots, \boldsymbol{\mu}_{ipM}] , \quad (1)$$

where $\boldsymbol{\mu}_{ipm}$, $m = 1, 2, \dots, M$ is the mean vector of the m th Gaussian component⁵. Similarly, a state cluster supervector $\mathbf{m}_i$ is created from the GMM of state cluster i .

STEP 10 : Collect the trigrapheme state supervectors $\{\mathbf{v}_{i1}, \mathbf{v}_{i2}, \dots, \mathbf{v}_{iN_i}\}$ as well as the state cluster supervector $\mathbf{m}_i$ of cluster i , and derive an eigenbasis from their correlation matrix using *weighted principal component analysis* (WPCA). The correlation matrix is computed as follows:

---

⁵Since the mixture weights are still tied among the trigrapheme states in a state cluster, the M Gaussian components in each state can be consistently ordered across all the member states in the cluster to create their supervectors.

$$\frac{1}{n_i} \sum_p n_{ip} (\hat{\mathbf{v}}_{ip} - \hat{\mathbf{m}}_i)(\hat{\mathbf{v}}_{ip} - \hat{\mathbf{m}}_i)' , \quad (2)$$

where $\hat{\mathbf{v}}_{ip}$ and $\hat{\mathbf{m}}_i$ are the standardized version of $\mathbf{v}_{ip}$ and $\mathbf{m}_i$ that are created by normalizing them with the diagonal covariance matrix; n_{ip} is the frame count of the trigrapheme state p in cluster i , and $n_i = \sum_p n_{ip}$ is the total frame counts of state cluster i . A comparison of using the standard PCA and WPCA in deriving the eigenbasis [22] shows that WPCA is more effective because the eigenvalues fall faster so that fewer eigenvectors are sufficient for representing the eigentrigrapheme space. Therefore, WPCA is used throughout this paper.

STEP 11 : Arrange the eigenvectors $\{\hat{\mathbf{e}}_{ik}, k = 1, 2, \dots, N_i\}$ in descending order of their eigenvalues λ_{ik} , and pick the top K_i (where $K_i \leq N_i$) eigenvectors to represent the eigenspace of state cluster i . These K_i eigenvectors are now called *eigentrigraphemes* of state cluster i . Note that different state clusters may have a different number of eigentrigraphemes.

2.3.2. Estimation of the Eigentrigrapheme Coefficients

After the derivation of the eigentrigraphemes, the supervector $\mathbf{v}_{ip}$ of any trigrapheme state p in cluster i is assumed to lie in the space spanned by the K_i eigentrigraphemes. Thus, we have

$$\mathbf{v}_{ip} = \mathbf{m}_i + \sum_{k=1}^{K_i} w_{ipk} \mathbf{e}_{ik} , \quad (3)$$

where $\mathbf{e}_{ik}, k = 1, 2, \dots, K_i$ is the rescaled version of the standardized eigenvector $\hat{\mathbf{e}}_{ik}$; $\mathbf{w}_{ip} = [w_{ip1}, w_{ip2}, \dots, w_{ipK_i}]$ is the eigentrigrapheme coefficient vector of trigrapheme state p in the trigrapheme state space of cluster i .

The eigentrigrapheme coefficient vector $\mathbf{w}_{ip}$ is estimated by maximizing the following penalized likelihood objective function $Q(\mathbf{w}_{ip})$:

$$Q(\mathbf{w}_{ip}) = L(\mathbf{w}_{ip}) - \beta R(\mathbf{w}_{ip}) , \quad (4)$$

where β is the regularization parameter that balances the dynamic ranges of the regularizer $R(\cdot)$ and the likelihood term $L(\cdot)$. $L(\mathbf{w}_{ip})$ is the likelihood of the training data of trigrapheme state p in cluster i , and is given by

$$L(\mathbf{w}_{ip}) = \text{constant} - \sum_{m,t} \gamma_{ipm}(t) (\mathbf{x}_t - \boldsymbol{\mu}_{ipm}(\mathbf{w}_{ip}))' C_{ipm}^{-1} (\mathbf{x}_t - \boldsymbol{\mu}_{ipm}(\mathbf{w}_{ip})) \quad (5)$$

where C_{ipm} and $\gamma_{ipm}(t)$ are the covariance and occupation probability of the m th Gaussian of trigrapheme state p in cluster i given observation $\mathbf{x}_t$.

To avoid the estimation of $\mathbf{w}_{ip}$ from over-fitting, the regularizer $R(\cdot)$ is introduced in the objective function. From Eqn.(3), one can see the $|\mathbf{w}_{ip}|$ is

the Euclidean distance of the trigrapheme state supervector $\mathbf{v}_{ip}$ from its state cluster supervector $\mathbf{m}_i$ in the trigrapheme state space. In [21], the following regularizer

$$R(\mathbf{w}_{ip}) = \sum_{k=1}^{K_i} \frac{w_{ipk}^2}{\lambda_{ik}} \quad (6)$$

is found effective. It has the following properties:

- The squared coefficient of each eigentrigrapheme, w_{ipk} , is inversely scaled by its eigenvalue so that a less informative eigentrigrapheme (with smaller eigenvalues) will have less influence on the “adapted” trigrapheme model. On the other hand, the trigrapheme can freely move along the more informative eigentrigraphemes (with larger eigenvalues) in the trigrapheme state space.
- Since the chosen regularizer automatically de-emphasizes the less informative eigentrigraphemes, we may avoid making a hard decision on the number of eigentrigraphemes K_i for each state cluster. Instead, we may simply use *all* eigentrigraphemes (by setting $K_i = N_i$) for every state cluster, eliminating the need to tune the value of K_i .
- When a trigrapheme state has a lot of training data, the likelihood term will become dominant in the objective function $Q(\mathbf{w}_{ip})$. As a result, the “adapted” model should converge to its Baum-Welch training estimate.
- On the other hand, for a trigrapheme state with limited amount of training data, the regularizer becomes dominant and it will try to make $|\mathbf{w}_{ip}|$ as small as possible. Thus, the “adapted” model will fallback to its state cluster GMM.

The above estimation method of the trigrapheme coefficients can be viewed as a penalized version of the *maximum-likelihood eigen-decomposition* (MLED) in eigenvoice adaptation [24], and it is called *penalized maximum-likelihood eigen-decomposition* (PMLED) in Fig. 1.

Differentiating the optimization function $Q(\mathbf{w}_{ip})$ of Eqn. (4) w.r.t. each eigentrigrapheme coefficient w_{ipk} , and setting each derivative to zero, we have,

$$\sum_{n=1}^{K_i} A_{ipkn} w_{ipn} + \beta \frac{w_{ipk}}{\lambda_{ik}} = B_{ipk}, \quad \forall k = 1, 2, \dots, K_i \quad (7)$$

where

$$\begin{aligned} A_{ipkn} &= \sum_m \mathbf{e}'_{ikm} C_{ipm}^{-1} \mathbf{e}_{inm} \left(\sum_t \gamma_{ipm}(t) \right) \\ B_{ipk} &= \sum_m \mathbf{e}'_{ikm} C_{ipm}^{-1} \left(\sum_t \gamma_{ipm}(t) (\mathbf{x}_t - \mathbf{m}_{im}) \right). \end{aligned}$$

The eigentrigrapheme coefficients may be easily found by solving the system of K_i linear equations represented by Eqn. (7), and the Gaussian mean of the m th mixture of trigrapheme state p in cluster i can be obtained from $\mathbf{v}_{ip}$ as

$$\boldsymbol{\mu}_{ipm} = \mathbf{m}_{im} + \sum_{k=1}^{K_i} w_{ipk} \mathbf{e}_{ikm} . \quad (8)$$

The eigentrigrapheme modeling procedure stops if either the estimation of eigentrigrapheme coefficients converges or the recognition accuracy of the trained models is maximum on a development data set. Otherwise, the training data are re-aligned with the current models, and the derivation of eigentrigraphemes and the estimation of their coefficients are repeated.

3. Experimental Evaluation

The effectiveness of our newly proposed eigentrigrapheme acoustic modeling method is evaluated on four under-resourced languages of South Africa with the assumption that no phonetic dictionaries are available. Since graphemes are the basic modeling units in grapheme-based modeling, word recognition accuracy is the main metric for the evaluation. Nonetheless, triphone-based systems were also built with the use of semi-automatically generated phonetic dictionaries so as to benchmark the results of our eigentrigrapheme results (where no dictionaries are used).

3.1. The Lwazi Speech corpus

The Lwazi project was set up to develop a telephone-based speech-driven information system to take advantage of the more and more popular use of telephones in South Africa nowadays. As part of the project, the Lwazi ASR corpus [27] was collected to provide the necessary speech and language resources in building ASR systems for all eleven official languages of South Africa.

The corpus was collected from approximately 200 speakers per language who are all first language speakers. Each speaker produced approximately 30 utterances, in which 16 of them are phonetically balanced read speech and the remainders are elicited short words such as answers to open questions, answers to yes/no questions, spelt words, dates, and numbers. All the data were recorded over a telephone channel and were transcribed only in words. Background noises, speaker noises, and partial words are marked in the orthographic transcriptions.

The Lwazi project also created a 5,000-word pronunciation dictionary for each language [28]. These dictionaries cover the most common words in the language but not all the words appearing in the corpus. Thus, for the phone-based experiments, the *DictionaryMaker* [29] software was used to generate dictionary entries for the words that are not covered by the Lwazi dictionaries. The given Lwazi dictionaries were used as the seed dictionaries⁶ for DictionaryMaker to

---

⁶For Afrikaans, the dictionary available at <http://sourceforge.net/projects/rcrl/files/AfrPronDict/> was used together with the Lwazi dictionary as the seed dictionary.

extract grapheme-to-phoneme conversion rules which were then applied to generate the phonetic transcriptions for the uncovered words for each language. The pronunciations suggested by DictionaryMaker were directly used without any modification.

Among the eleven South African official languages, four are chosen for this investigation. We looked at their ranks according to three different criteria:

- the human language technology (HLT) index [30]: the index indicates the total quantity of HLT activity for each language. Higher the index is, greater HLT development has been done.
- the phone recognition accuracy [31]: higher phone accuracy means a higher rank for the language.
- the amount of training data available [31]: language with more training data will be given a higher rank.

Table 1: Ranks of the four chosen South African languages in three aspects: their human language technology (HLT) indices, phone recognition accuracies, and amount of training data in the Lwazi corpus. (Smaller value implies a higher rank.)

Language	HLT Rank [30]	Phone Recognition [31]	Amount of Data [31]
Afrikaans	1	5	11
SA English	2	11	10
Sesotho	7	7	7
siSwati	9	3	1

Finally, the following four languages are chosen because they have a good mix of phone accuracies and HLT activities as shown in Table 1:

Afrikaans: Afrikaans is a Low Franconian, West Germanic language, originated from Dutch [32]. It has about 6 million native speakers and is the third largest language in South Africa. It is also spoken in South Africa’s neighbouring countries like Namibia, Botswana and Zimbabwe. It has relatively more resources [33], and more ASR related works [34, 35] have been done on it than other languages of South Africa. It is interesting to see that although Afrikaans has the least amount of training data in the corpus, its phone recognition result is quite good among the eleven South African languages.

South African (SA) English: SA English is the de facto South Africa’s lingua franca. It is spoken by about 3.6 million people in South Africa. SA English is evolved from British English but is highly influenced by Afrikaans and the other languages of the country.

Sesotho: Sesotho is a Southern Bantu language, closely related to other languages in the Sotho-Tswana language group. It has about 3.5 million native speakers and is the seventh largest language in South Africa.

siSwati: siSwati is also a Southern Bantu language, closely related to the Nguni language group. It has about 1 million native speakers and is the ninth largest language in South Africa.

Table 2: Information of the data sets of four South African languages used in this investigation. (OOV is *out-of-vocabulary*)

Data Set	#Speakers	#Utt.	Dur.(hr)	Vocab	OOV
Afrikaans Training	160	4,784	3.37	1,513	0.00%
Afrikaans Dev.	20	600	—	870	0.89%
Afrikaans Test	20	599	—	876	0.97%
SA English Training	156	4,665	3.98	1,988	0.00%
SA English Dev.	20	581	—	1,104	1.10%
SA English Test	20	597	—	1,169	1.68%
Sesotho Training	162	4,826	5.70	2,360	0.00%
Sesotho Dev.	20	600	—	1,096	1.86%
Sesotho Test	20	601	—	1,089	2.29%
siSwati Training	156	4,643	8.38	4,645	0.00%
siSwati Dev.	20	599	—	1,889	6.14%
siSwati Test	20	596	—	1,851	4.53%

Since the corpus does not define the training, development, and test set for each language, we did the partitions ourselves. The data sets used in our experiments are summarized in Table 2. It is interesting to see that languages with more training data (in terms of duration) have a higher percentage of out-of-vocabulary words in their test set.

3.2. Feature Extraction and Common Experimental Settings

The first 13 PLP coefficients, including c0, and their first and second order derivatives were used. These 39-dimensional feature vectors were extracted at every 10ms over a window of 25ms. Speaker-based cepstral mean subtraction and variance normalisation were performed.

The grapheme set and phone set of each language are the same as the ones defined in the Lwazi dictionaries. For all systems described below, transition probabilities of all triphones/trigraphemes of the same base phone/grapheme were tied together. Each triphone/trigrapheme model was a strictly left-to-right 3-state continuous-density hidden Markov model (CDHMM) with a Gaussian

Table 3: Perplexities of phone and word language models of the four South African languages.

Language	Data Set	Phone Perplexity	Word Bigram Perplexity
Afrikaans	Dev.	7.37 (trigram)	12.4
	Test	7.33 (trigram)	11.18
SA English	Dev.	7.50 (trigram)	13.28
	Test	7.76 (trigram)	11.18
Sesotho	Dev.	10.43 (bigram)	19.60
	Test	10.29 (bigram)	19.69
siSwati	Dev.	7.60 (trigram)	12.27
	Test	7.50 (trigram)	10.94

mixture density of at most $M = 16$ components per state. In addition, there were a 1-state short pause model and a 3-state silence model whose middle state was tied with the short pause state. Recognition was performed using the HTK toolkit [26] with a beam search threshold of 350. Only the annotated text data in the training set were used to train the corresponding language models. Both phone trigram language models and word bigram language models were estimated for the four languages except Sesotho, for which only phone bigrams could be reliably trained. Perplexities of the various language models on the development data and test data are shown in Table 3.

All system parameters such as the grammar factor, insertion penalty, regularization parameter β , number of GMM components M , number of tied states or state clusters, and so forth were optimized using the respective development data.

Table 4: Some system parameters of triphone modeling in the four South African languages.

Language	#Base Phones	#Cross-word Triphones	#Tied States in Conventional Models	#State Clusters in Eigentriphone Models
Afrikaans	37	5,203	617	332
SA English	44	7,167	988	362
Sesotho	41	4,061	741	624
siSwati	40	5,140	339	250

3.3. Phone and Word Recognition using Triphone-based HMMs

We first establish the triphone-based ASR benchmarks against which the trigrapheme-based models can be checked. Both conventional tied-state triphone HMM modeling and our new cluster-based eigentriphone modeling were

tried for the four under-resourced languages of South Africa. The number of base phones, the number of cross-word triphones in the training set, the optimal number of tied states in conventional HMM training, and the optimal number of state clusters in eigentriphone modeling for each language are summarized in Table 4.

Table 5: Phone recognition accuracy (%) of four South African languages. († The benchmark results in [31] used an older version of the Lwazi corpus and how the corpus were partitioned into training, development, and test sets is unknown.)

Language	Benchmark [31]†	Tied-state Triphone		Cluster-based Eigentriphone	
	Flat LM	Flat LM	N-gram LM	Flat LM	N-gram LM
Afrikaans	63.14	59.07	69.73 (trigram)	62.23	72.32 (trigram)
SA English	54.26	45.48	56.58 (trigram)	46.03	57.84 (trigram)
Sesotho	54.79	62.36	67.06 (bigram)	64.08	68.35 (bigram)
siSwati	64.46	64.76	71.45 (trigram)	68.19	74.13 (trigram)

3.3.1. Phone Recognition Results

Although word recognition accuracy will be the eventual evaluation metric for grapheme modeling, we also would like to report the phone recognition baselines of our triphone models for the sake of completeness. Phone recognition was performed on each of the four languages using no or flat LM as well as using its respective bigram/trigram LM. The results are given in Table 5. The following observations can be made:

- Our phone recognition results with flat LMs are quite different from those reported in [31]. There may be a few reasons:
  - To our knowledge, the Lwazi corpus has been evolving, and the corpus we obtained earlier this year is different from the older version used in [31].
  - Since there are no official test sets in the corpus, it is hard to compare recognition performance from different research groups.
  - Since the data are not manually labelled by professional transcribers, there is no ground truth which the results from different research groups can compare with.

Thus, it may not be meaningful to compare our phone recognition results with others. We believe it is good enough to see that our results are in the same ballpark as the others.

- SA English has substantially lower phone recognition accuracy: it is lower than that of the other three languages by more than 10% absolute. Although SA English has a few more phones in its phonetic inventory than

the other languages, and significantly more cross-word triphones to model (see Table 4), its phone trigram perplexity is actually similar to Afrikaans and siSwati. (Only bigram language model can be reliably estimated for Sesotho, and its value is expected to be higher than the phone trigram perplexity of the other three languages.) It means that the phone trigrams (as well as triphones) of SA English are more unevenly distributed in the training corpus.

The lower phone recognition accuracy of SA English may be simply due to its larger inventory of phones and triphones, making discrimination among them more difficult. Another plausible reason is that SA English is now the de facto lingua franca of South Africa. It is usually the language of choice for communication among people from different regions and ethnic groups of the country including immigrants from China and India. As a consequence, there are more allophonic variations in SA English, making it harder to recognize.

- The training speech data are not phonetically labelled by human transcribers. Instead, their phonetic transcriptions are generated semi-automatically by grapheme-to-phoneme conversion together with a small bootstrapping dictionary. From the big improvement of recognition performance when phone language models were used (vs. when no language models were used), we may conclude that phone language models trained from the generated phonetic transcriptions are good enough to improve phone recognition significantly.
- Triphone models estimated by our new cluster-based eigentriphone modeling method outperform triphone models estimated by conventional tied-state HMM training by an average of 6.19% relative over the four languages.

Table 6: Word recognition accuracy (%) of four South African languages.

Language	Tied-state Triphone		Cluster-based Eigen-trigrapheme/-triphone	
	Trigrapheme	Triphone	Trigrapheme	Triphone
Afrikaans	89.39	89.73	89.87	90.73
SA English	78.30	83.12	79.57	83.72
Sesotho	75.67	75.57	76.35	76.77
siSwati	80.04	79.79	80.67	80.29

3.3.2. Word Recognition Results

The word recognition performance of the triphone-based systems are shown in Table 6. We can see that

- With no surprise, Sesotho, having the highest LM perplexity (see Table 3), has the lowest recognition accuracy.
- For the other three languages, namely Afrikaans, SA English, and siSwati, which all have similar word bigram perplexity, their word recognition performance is well correlated with their vocabulary size and OOV figure. Afrikaans has the best word recognition accuracy, and yet there are only 1,513 words in its vocabulary with 0.97% OOV. On the other hand, siSwati has the worst performance, and its vocabulary size is 4,645 with 4.53% OOV, which are 3–4 times of those of Afrikaans (see Table 4).
- Although SA English has the poorest phone recognition accuracy, its word recognition performance is second among the four languages. It does not only show the limitation of using phone recognition accuracy to predict word recognition performance, but also the effectiveness of a good n-gram language model for word recognition.
- Cluster-based eigentriphone modeling outperforms conventional tied-state HMM training by an average of 5.17% relative over the four languages.

Table 7: Some system parameters used in trigrapheme modeling of the four South African languages. (The numbers of possible base graphemes are 43, 26, 27, 26 for the four languages but not all of them are seen in the corpus.)

Language	#Seen Base Graphemes	#Cross-word Trigraphemes	#Tied States in Conventional Models	#State Clusters in Eigentriphrase Models
Afrikaans	31	3,458	728	332
SA English	26	4,125	1,630	547
Sesotho	25	3,072	543	543
siSwati	25	3,826	392	255

3.4. Word Recognition using Trigrapheme-based HMMs

Similar acoustic models were developed using trigraphemes; there is no need for a phonetic dictionary in the process. The number of base graphemes actually observed in the corpus, the number of cross-word trigraphemes in the training set, the optimal number of tied states in conventional HMM training, and the optimal number of state clusters in eigentriphrase modeling for each language are summarized in Table 7. The word recognition results of the various trigrapheme-based systems are shown in Table 6 together with the results from the corresponding triphone-based systems so that they can be easily compared.

Besides the observations that are mentioned in triphone-based systems in Section 3.3.2, the following additional observations are well noted.

- Except for SA English, our trigrapheme-based systems performs basically the same as their triphone-based counterparts even without the knowledge of a phonetic dictionary. In fact, trigrapheme-based systems even outperform their triphone-based counterpart in siSwati though insignificantly. The results suggest that there is a consistent mapping between the pronunciation of Afrikaans, Sesotho, and siSwati and their graphemes.
- Trigrapheme-based systems perform much worse than triphone-based systems in SA English. This is expected. Similar results have been reported for English [17]. Besides the reason mentioned in the Introduction Section that the pronunciation of English has developed away from its written form over time, the particularly great allophonic variations in SA English (which is also reflected in its phone recognition accuracy) further compromise the word recognition effort.
- Once again, our new cluster-based eigentrigrapheme modeling consistently performs better than conventional tied-state trigrapheme HMM training. It has an average gain of 4.08% relative over the four languages.

4. Conclusions

Most state-of-the-art automatic speech recognition (ASR) systems are developed using phonetic acoustic models. However, for many developing or under-developed countries in the world, the adoption of human language technologies has been dragged down by the lack of speech and language resources, which are usually costly and take a lot of human expertise to acquire. Graphemic acoustic modeling mitigates the problem as it does not require a phonetic dictionary. In this paper, we port a new acoustic modeling method called cluster-based eigentriphone modeling which has been shown to outperform the conventional tied-state triphone HMM training in phone-based ASR systems to grapheme-based ASR for under-resourced languages. We call the new method *cluster-based eigentrigrapheme acoustic modeling*.

For four under-resourced languages of South Africa (SA), namely, Afrikaans, SA English, Sesotho, and siSwati, it is shown that in terms of word recognition performance, trigrapheme-based ASR is as good as triphone-based ASR with the exception of SA English. The worse performance of trigrapheme-based ASR on SA English is not unexpected, since SA English is a variation of British English, and it is well known that grapheme-based ASR does not perform well on the latter. In particular, trigrapheme acoustic models trained by our new eigentrigrapheme modeling method consistently outperforms the trigrapheme models trained by conventional tied-state HMM training, achieving a relative reduction in the word error rates of the four SA languages by an average of 4.08%. Trigrapheme HMM states trained by the eigentrigrapheme modeling method are distinct from each other — the quantization error among the member states of a tied state in conventional HMM is avoided — and should be more discriminative.

In the future, we would like to investigate other cluster definitions for under-resourced language ASR, especially when the amount of acoustic training data is even smaller. The effect of discriminative training on the inherently distinctive models produced by eigentrigrapheme modeling will be studied as well.

5. Acknowledgments

This research is partially supported by the Research Grants Council of the Hong Kong SAR under the grant numbers FSGRF12EG31, FSGRF13EG20, and SRFI11EG15.

References

- [1] K. U. Ogbureke, J. Carson-Berndsen, Framework for cross-language automatic phonetic segmentation, in: Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, 2010.
- [2] V.-B. Le, L. Besacier, Automatic speech recognition for under-resourced languages: Application to Vietnamese language, IEEE Transactions on Audio, Speech and Language Processing 17 (2009) 1471–1482.
- [3] J. Kohler, Multi-lingual phoneme recognition exploiting acoustic-phonetic similarities of sounds, in: Proceedings of the International Conference on Spoken Language Processing, 1996.
- [4] P. Beyerlein, et al., Towards language independent acoustic modeling, in: Proceedings of the IEEE Automatic Speech Recognition and Understanding Workshop, 1999.
- [5] J. J. Sooful, E. C. Botha, An acoustic distance measure for automatic cross-language phoneme mapping, in: Proceedings of Pattern Recognition Association of South Africa, 2001.
- [6] D. Povey, et al., Subspace Gaussian mixture models for speech recognition, in: Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, 2010.
- [7] L. Burget, P. Schwarz, M. Agarwal, P. Akyazi, K. Feng, A. Ghoshal, O. Glembek, N. Goel, M. Karafit, D. Povey, A. Rastrow, R. Rose, S. Thomas, Multilingual acoustic modeling for speech recognition based on subspace Gaussian mixture models, in: Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, 2010.
- [8] L. Lu, A. Ghoshal, S. Renals, Regularized subspace gaussian mixture models for cross-lingual speech recognition, in: Automatic Speech Recognition and Understanding (ASRU), 2011 IEEE Workshop on, 2011, pp. 365–370.

- [9] H. Meng, S. Hunnicutt, S. Seneff, V. Zue, Reversible letter-to-sound generation based on parsing word morphology, *Speech Communications* 18 (1996) 47–63.
- [10] J. R. Bellegarda, Unsupervised, language-independent grapheme-to-phoneme conversion by latent analogy, in: *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2003.
- [11] M. Davel, E. Barnard, The efficient creation of pronunciation dictionaries: human factors in bootstrapping, in: *Proceedings of Interspeech*, 2004.
- [12] O. Andersen, R. Kuhn, A. Lazarides, P. Dalsgaard, J. Haas, E. Noth, Comparison of two tree-structured approaches for grapheme-to-phoneme conversion, in: *Proceedings of the International Conference on Spoken Language Processing*, 1996.
- [13] E. G. Schukat-Talamazzini, H. Niemann, W. Eckert, T. Kuhn, S. Rieck, Automatic speech recognition without phonemes, in: *Proceedings of the European Conference on Speech Communication and Technology*, 1993.
- [14] S. Kanthak, H. Ney, Context-dependent acoustic modeling using graphemes for large vocabulary speech recognition, in: *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2002.
- [15] S. H. P. Charoenpornasawat, T. Schultz, Thai grapheme-based speech recognition, in: *In Proceedings of the Human Language Technology Conference of the North American Chapter of the ACL*, 2006.
- [16] P. T. Daniels, W. Bright (Eds.), *The World’s Writing Systems*, Oxford University Press, 198 Madison Avenue, New York, NY 10016, USA, 1996.
- [17] S. Stuker, Acoustic modelling for under-resourced languages, Ph.D. thesis, Universität Fridericiana zu Karlsruhe (TH) (2009).
- [18] S. Stuker, Modified polyphone decision tree specialization for porting multilingual grapheme based ASR systems to new languages, in: *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2008.
- [19] S. Kanthak, H. Ney, Multilingual acoustic modeling using graphemes, in: *Proceedings of the European Conference on Speech Communication and Technology*, 2003.
- [20] T. Ko, B. Mak, Eigentriphones: A basis for context-dependent acoustic modeling, in: *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2011.
- [21] T. Ko, B. Mak, A fully automated derivation of state-based eigentriphones for triphone modeling with no tied states using regularization, in: *Proceedings of Interspeech*, 2011.

- [22] T. Ko, B. Mak, Derivation of eigentriphones by weighted principal component analysis, in: Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, 2012.
- [23] T. Ko, B. Mak, Eigentriphones for context-dependent acoustic modeling, submitted to IEEE Transactions on Audio, Speech and Language Processing (2012).
- [24] R. Kuhn, J.-C. Junqua, P. Nguyen, N. Niedzielski, Rapid speaker adaptation in eigenvoice space, IEEE Transactions on Speech and Audio Processing 8 (2000) 695–707.
- [25] S. J. Young, J. J. Odell, P. C. Woodland, Tree-based state tying for high accuracy acoustic modelling, in: Proceedings of the Workshop on Human Language Technology, 1994.
- [26] S. Young, et al., The HTK Book (Version 3.4), University of Cambridge, 2006.
- [27] Meraka-Institute, Lwazi ASR corpus, Online: <http://www.meraka.org.za/lwazi>, 2009.
- [28] M. Davel, O. Martirosian, Pronunciation dictionary development in resource-scarce environments, in: Proceedings of Interspeech, 2009.
- [29] M. Tempest, M. Davel, DictionaryMaker 2.16 user manual, <http://dictionarymaker.sourceforge.net.>, 2009.
- [30] A. Sharma-Grover, G. B. van Huyssteen, M. W. Pretorius, An HLT profile of the official South African languages, in: Proceedings of the Second Workshop on African Language Technology (AfLaT 2010), 2010.
- [31] C. van Heerden, E. Barnard, M. Davel, Basic speech recognition for spoken dialogues, in: Proceedings of Interspeech, 2009.
- [32] G. B. van Huyssteen, S. Pilon, Rule-based conversion of closely-related languages: A Dutch-to-Afrikaans convertor, in: Proceedings of the 20th Annual Symposium of the Pattern Recognition Association of South Africa, 2009.
- [33] J. C. Roux, P. H. Louw, T. R. Niesler, The African speech technology project: An assessment, in: Proceedings of LREC, 2004.
- [34] F. de Wet, A. de Waal, G. B. van Huyssteen, Developing a broadband automatic speech recognition system for Afrikaans, in: Proceedings of Interspeech, 2011.
- [35] H. Kamper, T. Niesler, Multi-accent speech recognition of Afrikaans, black and white varieties of South African English, in: Proceedings of Interspeech, 2011.