Articulatory Feature based Continuous Speech Recognition using Probabilistic Lexical Modeling

Ramya Rasipuram^{a,b,*}, Mathew Magimai.-Doss^a

{ramya.rasipuram,mathew}@idiap.ch

^aIdiap Research Institute, Martigny, Switzerland

^bEcole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Abstract

Phonological studies suggest that the typical subword units such as phones or phonemes used in automatic speech recognition systems can be decomposed into a set of features based on the articulators used to produce the sound. Most of the current approaches to integrate articulatory feature (AF) representations into an automatic speech recognition (ASR) system are based on a deterministic knowledge-based phoneme-to-AF relationship. In this paper, we propose a novel two stage approach in the framework of probabilistic lexical modeling to integrate AF representations into an ASR system. In the first stage, the relationship between acoustic feature observations and various AFs is modeled. In the second stage, a probabilistic relationship between subword units and AFs is learned using transcribed speech data. Our studies on a continuous speech recognition task show that the proposed approach effectively integrates AFs into an ASR system. Furthermore, the studies show that either phonemes or graphemes can be used as subword units. Analysis of the probabilistic relationship captured by the parameters has shown that the approach is capable of adapting the knowledge-based phoneme-to-AF representations using speech data; and allows different AFs to evolve asynchronously.

Keywords: Automatic speech recognition; articulatory features; probabilistic lexical modeling, Kullback-Leibler divergence based hidden Markov model; phoneme subword units; grapheme subword units.

1. Introduction

Articulatory features describe the properties of speech production, i.e., each sound unit of a language, a phone or a phoneme, can be decomposed into a set of features based on the articulators used to produce it. The use of articulatory

*Corresponding author

feature (AF) representations in an automatic speech recognition (ASR) system is motivated by their abilities such as:

- Better pronunciation modeling: AFs are hypothesized to capture acoustic variation at finer level than the phoneme-based representation (Deng et al., 1997; Richardson et al., 2003; Livescu et al., 2008)
- Robustness to noise: Different AFs may have variable noise sensitivity. The “divide and conquer” approach provides a framework to exploit the variable noise sensitivity of AFs (Kirchhoff et al., 2002).
- Multilingual and crosslingual portability: AFs can provide better sharing capabilities than phonemes across languages (Stüker et al., 2003; Lal and King, 2013; Siniscalchi et al., 2012).

To incorporate the articulatory knowledge in an ASR system, the following main concerns have to be addressed:

1. AF representations: There exist different types of articulatory representations of speech, e.g., binary features, multivalued features, and government phonological features. AFs defined by Chomsky and Halle (1968) are binary valued features, for example +voice and -voice, +sonorant and -sonorant. However, according to Ladefoged (1993), it is more natural to allow AFs to take multiple values. In government phonological feature system, speech sounds are destructured into a set of primes and can be represented by fusing them structurally (Harris, 1994). In the paper, we have used the multivalued AFs, because it has been argued that they better represent non-binary parameters such as height of vowels and place of articulation (Ladefoged, 1993).
2. Estimation of AFs from acoustic speech signal: In the literature, many approaches have been explored to extract AFs from the acoustic speech signal. For example, techniques based on acoustic-to-articulatory feature codebooks (Hogden et al., 1996; Suzuki et al., 1998), artificial neural networks (Livescu et al., 2008; Kirchhoff et al., 2002; Chang, 2002; Rasipuram and Magimai-Doss, 2011b), support vector machines (Juneja and Espy-Wilson, 2004; Scharenborg et al., 2007), Gaussian mixture models (Metze and Waibel, 2002; Stüker et al., 2003), hidden Markov models (Hiroya and Honda, 2004), conditional random fields (Prabhavalkar et al., 2011) nearest neighbour (Næss et al., 2011), dynamic Bayesian networks (Frankel and King, 2005; Frankel et al., 2007) are used.
3. Integration: Integrating AFs into the conventional hidden Markov model (HMM) based ASR framework is a challenging task mainly because of the multiple AF estimators. The dynamic Bayesian network (DBN) based approaches for AF integration preserve the articulatory representation in DBN state space (Livescu and Glass, 2004; Livescu et al., 2008; King et al.,

2007). These approaches have shown promising results in lexical access¹ experiments. Posterior probabilities of AFs can be transformed for use as features in tandem speech recognition systems (Cetin et al., 2007a,b; Lal and King, 2013). Posterior probabilities of AFs are also used to enhance phoneme-based acoustic models (Cetin et al., 2007a,b; Lal and King, 2013). These approaches however lose other benefits of articulatory representation such as finer granularity and asynchronous evolution.

In this paper, we propose an approach in the framework of probabilistic lexical modeling to integrate multivalued AFs. In a probabilistic lexical model based ASR system, the relationship between subword units in the lexicon and acoustic feature observations is factored into two models using latent variables: An acoustic model which models the relationship between acoustic feature observations and the latent variables; and a lexical model which models a probabilistic relationship between the subword units in the lexicon and the latent variables. In this paper, we show that by choosing the latent variables as multiple multivalued AFs, the approach effectively integrates AFs into the HMM-based ASR framework. The lexical model parameters in the proposed approach capture a probabilistic relationship between subword units and AFs learned through transcribed speech data.

The potential of the proposed approach for AF integration is demonstrated on a continuous speech recognition task through experiments and comparisons with the tandem approach. In the proposed framework we explore the use of domain-independent data for acoustic model training; and phonemes and graphemes as subword units. Furthermore, through the analysis of the lexical model parameters we show that the approach adapts the knowledge-based phoneme-to-AF or grapheme-to-AF relationship and allows different AFs to evolve asynchronously.

The rest of the paper is organized as follows: Section 2 gives an overview of the HMM-based ASR and the framework of probabilistic lexical modeling. Section 3 presents the literature review of approaches that integrate multivalued AFs for ASR in the light of the background information given in Section 2. In Section 4, the approach for AF integration is presented and the contributions of the present paper with respect to prior work are elaborated. Sections 5 and 6 present the experimental setup and results, respectively. Section 7 presents an analysis on the subword-unit-to-AF relationship captured by the lexical model parameters. Finally, in Section 8 we provide a discussion and conclusion.

2. Background

The goal of ASR is to find the most likely word sequence $W^* = [w_1, \dots, w_m, \dots, w_M]$ given the acoustic observation sequence $X =$

¹The task of lexical access involves predicting a word given its phonetic or broad phonetic transcription (Huttenlocher and Zue, 1984).

$[\mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T]$ where M is the number of words in the utterance and T represents the number of frames in the speech signal. The most likely word sequence W^* given the acoustic observation sequence is obtained as follows:

$$W^* = \arg \max_{W \in \mathcal{W}} P(W|X) \quad (1)$$

$$= \arg \max_{W \in \mathcal{W}} p(X|W)P(W) \quad (2)$$

where $\mathcal{W}$ denotes the set of all possible word sequences. The first term on the right hand side of Eqn (2) denotes the acoustic likelihood and the second term denotes the language model probability.

In general, speech recognition systems model words as a sequence of subword units, which are further modeled as a sequence of HMM states. The sequence of subword units for a word is given by its pronunciation model as specified in the pronunciation lexicon. The acoustic likelihood in an HMM-based ASR system is computed as follows:

$$p(X|W, \Theta_A) = \sum_{Q \in \mathcal{Q}} p(X|Q, W, \Theta_A)P(Q|W, \Theta_A) \quad (3)$$

$$= \sum_{Q \in \mathcal{Q}} p(X|Q, \Theta_A)P(Q|W, \Theta_A) \quad (4)$$

$$\approx \max_{Q \in \mathcal{Q}} p(X|Q, \Theta_A)P(Q|W, \Theta_A) \quad (5)$$

$$\approx \max_{Q \in \mathcal{Q}} \left[\prod_{t=1}^T p(\mathbf{x}_t|q_t, \Theta_A)P(q_t|q_{t-1}, \Theta_A) \right] \quad (6)$$

The set Θ_A includes the parameters of the acoustic likelihood estimator. In Eqn (3), the acoustic likelihood is obtained by summing over all possible state sequences $\mathcal{Q}$ where each $Q = [q_1, \dots, q_t, \dots, q_T]$ denotes a sequence of HMM states corresponding to a word sequence hypothesis. Eqn (4) assumes that acoustic likelihood is independent of words given the state sequence. In Eqn (5), a Viterbi approximation is employed where the sum over all possible state sequences is replaced with the most probable state sequence. Eqn (6) arises from the two HMM assumptions, i.e., acoustic feature observations are conditionally independent of each other and the HMM state at time t depends only on the HMM state at time $t - 1$.

In subword unit based ASR systems, HMM states represent lexical units, i.e., $q_t \in \mathcal{L} = \{l^1, \dots, l^i, \dots, l^I\}$ and I is the number of lexical units. If context-independent phonemes are used as subword units then the number of lexical units $I = M \times K$ where K is the number of context-independent subword units in the lexicon and M is the number of HMM states for each context-independent phoneme. If context-dependent phonemes are used as subword units then the number of lexical units $I = M \cdot K^{c_r+c_l+1}$ where c_l is the preceding context length, c_r is the following context length. Typically, each context-independent or context-dependent phoneme is modeled with three HMM states, i.e., $M = 3$.

In HMM-based ASR systems, the relationship between lexical units and acoustic features is not always modeled directly. As we will see in the remainder of the section, the relationship is typically modeled through intermediate acoustic units (for example, clustered context-dependent subword states).

2.1. Framework of Probabilistic Lexical Modeling

In the framework of probabilistic lexical modeling (Rasipuram and Magimai-Doss, 2015), relationship between the acoustic feature observation $\mathbf{x}_t$ and the lexical unit l^i is factored through a *latent* variable a^d as follows:

$$p(\mathbf{x}_t | q_t = l^i, \Theta_A) = \sum_{d=1}^D p(\mathbf{x}_t, a^d | q_t = l^i, \Theta_A) \quad (7)$$

$$= \sum_{d=1}^D p(\mathbf{x}_t | a^d, q_t = l^i, \theta_a, \theta_l) \cdot P(a^d | q_t = l^i, \theta_l) \quad (8)$$

$$= \sum_{d=1}^D \underbrace{p(\mathbf{x}_t | a^d, \theta_a)}_{\text{acoustic model}} \cdot \underbrace{P(a^d | q_t = l^i, \theta_l)}_{\text{lexical model}} \quad (9)$$

The parameters of the acoustic likelihood estimator Θ_A encompass the *acoustic model* (θ_a), the *pronunciation lexicon* (θ_{pr}) and the *lexical model* (θ_l) parameters, therefore, $\Theta_A = \{\theta_a, \theta_{pr}, \theta_l\}$. The relationship in Eqn (9) is as a result of the assumption that given a^d , $p(\mathbf{x}_t | a^d, q_t = l^i, \theta_a, \theta_l)$ is independent of $q_t = l^i$. In Eqn (9), $p(\mathbf{x}_t | a^d, \theta_a)$ is the acoustic unit likelihood and $P(a^d | q_t = l^i, \theta_l)$ is the probability of the latent variable given the lexical unit. We refer to $p(\mathbf{x}_t | a^d, \theta_a)$ as the acoustic model, $P(a^d | q_t = l^i, \theta_l)$ as the lexical model, the latent variable a^d as the acoustic unit, the set of acoustic units $\mathcal{A} = \{a^1, \dots, a^d, \dots, a^D\}$ and D as the number of acoustic units.

Figure 1 shows the Bayesian network of an ASR system that uses the factorization of Eqn (9). The lexical unit is given deterministically by the current word and its subword units. The lexical unit is mapped to all the acoustic units and the acoustic feature observation is conditioned on all the acoustic units.

Figure 1: The graphical model representation of a system incorporating probabilistic lexical modeling. The D in the middle of the node indicates that there are D acoustic units.

2.2. Lexical and Acoustic Units

In the case of context-independent ASR systems, the lexical unit set $\mathcal{L}$ and the acoustic unit set $\mathcal{A}$ are knowledge driven and defined based on the subword units in the pronunciation lexicon. The number of lexical units or acoustic units $I = D = M \times K$, typically, $M = 3$.

In the case of context-dependent ASR systems, the number of lexical units $I = M \cdot K^{c_r + c_l + 1}$. Generally, not all context-dependent subword units will appear sufficiently often in the training data. Hence a sharing approach is used to enable multiple lexical units to share an acoustic model. This is done using a decision-tree based state clustering and tying technique that uses a pronunciation lexicon, linguistic knowledge to prepare a phonetic question set and acoustic data (Young et al., 1994). The number of acoustic units D varies depending on hyper parameters such as the state occupancy count and the log-likelihood threshold that are used during decision-tree based state clustering. However, the number of acoustic units D is well below the number of lexical units I . The resulting acoustic units are typically referred as clustered context-dependent states or tied-HMM states.

Other possibilities for the choice of the acoustic units are fenones (Bahl et al., 1988), senones (Hwang and Huang, 1992), automatically derived units from the acoustic data (Holter and Svendsen, 1997) etc. In this paper, we show that HMM-based ASR systems can be built using multivalued AFs as the acoustic units.

2.3. Acoustic Model

The acoustic model which models the relationship between the acoustic feature observation $\mathbf{x}_t$ and the acoustic unit a^d can be trained using either generative approaches such as Gaussian mixture models (GMMs) or discriminative approaches such as artificial neural networks (ANN).

GMMs (Rabiner, 1989): In the GMM approach, the acoustic model score $p(\mathbf{x}_t|a^d, \theta_a)$ is estimated given a mixture of C^d Gaussians that model an acoustic unit a^d , i.e.,

$$p(\mathbf{x}_t|a^d, \theta_a) = \sum_{c=1}^{C^d} w_c^d \mathcal{N}(\mathbf{x}_t, \mu_c^d, \Sigma_c^d) \quad (10)$$

where w_c^d , μ_c^d and Σ_c^d are the weight, means and covariances of the mixture component c of the acoustic unit a^d . In the literature, there are various variants of the GMM approach such as semi-continuous GMMs (Huang and Jack, 1989), subspace GMMs (Povey et al., 2011a).

ANN (Morgan and Bourlard, 1995): The ANN computes the probability of acoustic units given the acoustic feature observations $P(a^d|\mathbf{x}_t, \theta_a)$ which is then converted to scaled-likelihood, i.e.,

$$p_{sl}(\mathbf{x}_t|a^d, \theta_a) = \frac{p(\mathbf{x}_t|a^d, \theta_a)}{p(\mathbf{x}_t)} = \frac{P(a^d|\mathbf{x}_t, \theta_a)}{P(a^d)} \quad (11)$$

2.4. Lexical Model

The lexical model in an ASR system can be deterministic or probabilistic.

2.4.1. Deterministic Lexical Model based ASR Approaches

When the lexical model is deterministic, each lexical unit l^i is deterministically mapped to an acoustic unit a^j ($l^i \mapsto a^j$), i.e.,

$$P(a^d | q_t = l^i, \theta_l) = \begin{cases} 1, & \text{if } d = j; \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

As a result of the deterministic mapping, the only term contributing to the summation in Eqn (9) is the acoustic unit that is mapped to a lexical unit.

In standard HMM-based ASR approaches such as HMM/GMM (Rabiner, 1989) and hybrid HMM/ANN (Morgan and Bourlard, 1995), the lexical model is deterministic. In the case of context-independent ASR systems, it is a *knowledge-based look-up table* that maps each lexical unit to an acoustic unit. In the case of context-dependent ASR systems, typically the lexical model is obtained through clustering and state tying using *decision trees*. The decision trees map each context-dependent subword unit to a tied HMM state (or an acoustic unit). Figure 2 illustrates various steps in an HMM-based ASR system where the lexical model is deterministic. In this illustration, each context-dependent subword unit is composed of one lexical unit. However, normally each context-dependent subword unit is represented with three lexical units.

Figure 2: Sequence of steps taken in a context-dependent HMM-based ASR system where the lexical model is deterministic.

2.4.2. Probabilistic Lexical Model based ASR Approaches

The two conditions, namely, $0 < P(a^d | l^i, \theta_l) < 1$ and $\sum_{d=1}^D P(a^d | l^i, \theta_l) = 1$ in Eqn (9) characterize an ASR approach where each lexical unit is probabilistically related to all acoustic units. In (Rasipuram and Magimai-Doss, 2015), we showed that there are approaches such as probabilistic classification of HMM states (Luo and Jelinek, 1999), tied posteriors (Rottland and Rigoll, 2000) and Kullback-Leibler divergence based hidden Markov model (KL-HMM) (Aradilla et al., 2008) where the relationship between lexical units and acoustic units is probabilistic. The experimental studies in our previous work indicated that

the KL-HMM approach performs better than that of the tied posterior approach on various ASR tasks (Rasipuram and Magimai-Doss, 2015). Therefore, in this paper, we use the KL-HMM approach for probabilistic lexical modeling. The KL-HMM approach is briefly explained in Section 4.1.

2.5. Advantages of Probabilistic Lexical Modeling

Pronunciation variability modeling: Standard HMM-based ASR systems like HMM/GMM and hybrid HMM/ANN deterministically model the relationship between lexical units and acoustic feature observations. As a result of the deterministic relationship, these systems rely on a well developed phoneme lexicon to handle the variability in the acoustic training data. However, when the pronunciations in the lexicon do not reflect the underlying speech data then such a model may poorly represent the training data. For example, this can happen in the case of non-native speakers (when pronunciations normally reflect native speakers) or in the case of spontaneous and conversational speech (when spoken words are pronounced differently from lexicon pronunciations) or in the case of a grapheme lexicon (where pronunciations are based on the orthography of the word). To account for such a variation, typically, phoneme-based ASR systems add pronunciation variants to the lexicon (Strik and Cucchiaroni, 1999). However, the manual addition of pronunciation variants may require explicit human knowledge. In the context of pronunciation variability modeling, it has been shown that the limitation of the standard HMM/GMM system imposed by deterministic mapping can be handled by modeling a probabilistic relationship between lexical and acoustic units (Hain and Woodland, 1999; Saraclar et al., 2000; Hain, 2005; Rasipuram and Magimai-Doss, 2013b, 2015).

Resource optimization: In probabilistic lexical model based approaches such as KL-HMM and tied posteriors, the acoustic model and the lexical model are trained one after another and can be trained on an independent set of resources. For example, the acoustic model can be trained on resources from resource-rich languages and domains whereas the lexical model can be trained on a relatively small amount of target language data (Imseng et al., 2012; Rasipuram and Magimai-Doss, 2015).

Flexibility in the choice of acoustic and lexical units: In probabilistic lexical model based ASR approaches, it is not necessary that the subword unit set used for defining the acoustic units should be the same as the subword unit set used for defining the lexical units. The lexical model can capture the relationship between the distinct subword unit sets through acoustics. This flexibility has been exploited to build ASR systems where the acoustic unit set is based on phonemes and the lexical unit set is based on graphemes (Magimai-Doss et al., 2011; Rasipuram and Magimai-Doss, 2013a). Furthermore, lexical and acoustic units can model different contextual units. For instance, lexical units can be based on context-dependent subword units while the acoustic units can be based on context-independent subword units (Rottland and Rigoll, 2000; Magimai-Doss et al., 2011; Imseng et al., 2012).

In Section 4.2, we propose a novel approach to integrate articulatory features into HMM-based ASR in the framework of probabilistic lexical modeling

that can exploit all the above advantages.

3. Literature Survey

There has been a sustained interest in incorporating speech production knowledge into an ASR system for reasons already stated in Section 1. AFs have been incorporated at various levels of an ASR system. Here we provide a brief overview of ASR systems that used multivalued AFs according to the background of the previous section.

3.1. *Lexical units are phonemes and acoustic units are AFs*

In these works, the acoustic units are based on AFs or both AFs and phonemes (Metze and Waibel, 2002; Stüker et al., 2003; Juneja and Espy-Wilson, 2004; Livescu et al., 2008). Each acoustic unit is modeled with a GMM or with discriminative classifiers like ANNs or support vector machines. The lexical model is deterministic, i.e., each phoneme-based lexical unit is deterministically mapped to its AF attributes. The scores from different AF-based acoustic models are combined to arrive at the local emission score $p(x_t|q_t = l^i)$. On continuous speech recognition and cross-lingual adaptation tasks, the use of AF-based acoustic models in combination with phoneme-based acoustic models resulted in a relative reduction in word error rate (WER) of about 5-10% compared to the use of phoneme-based acoustic models alone (Metze and Waibel, 2002; Stüker et al., 2003).

3.2. *Lexical units are AFs and acoustic units are AFs*

These are the systems analogous to standard HMM-based ASR systems where both lexical and acoustic units are either based on context-independent or context-dependent subword units. However, unlike standard ASR systems, the subword units are AFs determined from the AF-based pronunciation lexicon (Deng et al., 1997; Richardson et al., 2003; Kirchhoff, 1996; Wester et al., 2004; Livescu et al., 2008). The AF-based pronunciation lexicon transcribes each word in terms of the positions of the articulators. Each AF is associated with its own hidden state variable. The multiple hidden state variables can follow an independent path to a certain extent and can allow certain amount of asynchrony. In the initial works, hidden state variables of various AFs were required to re-synchronize at phoneme level (Deng et al., 1997; Richardson et al., 2003) or syllable level (Kirchhoff, 1996). In more recent works, the flexible DBN framework allows synchronization to happen at word level or even across word boundaries (Livescu and Glass, 2004; Livescu et al., 2008). When the lexical units are based on context-dependent AFs, the acoustic units are typically clustered context-dependent subword states obtained using decision tree-based state tying methods. These systems have obtained improvements in lexical access experiments (Livescu and Glass, 2004; Livescu et al., 2008; Jyothi et al., 2011).

1
2
3
4
5
6
7
8
9 An approach was proposed by Jyothi et al. (2013) to convert a DBN-based
10 pronunciation model into an equivalent set of factored WFSTs. The utility of
11 this approach was demonstrated using a phoneme-based pronunciation model
12 on isolated word and continuous word speech recognition tasks; and using an
13 AF-based pronunciation model on lexical access tasks. Along the similar lines,
14 an approach was outlined to convert an AF-based DBN pronunciation model
15 into equivalent WFSTs for ASR (Jyothi, 2013, Chapter 8).
16

17 3.3. *Lexical units are phonemes and acoustic units are phonemes*

18 In these systems, similar to standard HMM-based ASR systems, both lexi-
19 cal and acoustic units are based on phonemes. However, AF representations
20 are used as auxiliary information to enhance the performance of the acous-
21 tic model (Kirchhoff et al., 2002; Siniscalchi et al., 2012). For example, the
22 acoustic model can be seen as two stage classifier. In the first stage, a set of
23 AF-based ANNs model the relationship between acoustic features and AFs. In
24 the second stage, a phoneme-based ANN models the relationship between all
25 AFs and phonemes. The resulting phoneme-based ANN is used as an acoustic
26 model in hybrid HMM/ANN systems. These systems have achieved a relative
27 reduction in WER of about 5-6% on noise robust ASR tasks and cross-lingual
28 ASR tasks compared to the systems where acoustic-to-phoneme information is
29 directly modeled (Kirchhoff et al., 2002; Siniscalchi et al., 2012).
30

31 Alternatively, AF-based neural networks have also been used in tandem
32 speech recognition systems (Cetin et al., 2007a,b; Livescu et al., 2008; Lal
33 and King, 2013). In the tandem approach, the posterior probabilities of AFs
34 and/or phonemes replace conventional cepstral features in HMM-based ASR
35 systems (Hermansky et al., 2000). In order to model the output of an ANN
36 that is typically non-Gaussian with GMMs, posterior probabilities of the acous-
37 tic units are Gaussianized using the log function and then decorrelated using
38 the Karhunen-Loeve transform (KLT).
39

40 For tandem systems, the use of AF-based ANNs trained on language-
41 independent data was investigated and compared with the use of AF-based
42 ANNs trained on language-dependent data (Cetin et al., 2007b). Cetin et al.
43 (2007b) observed that the AF-based ANNs trained on language-independent
44 data reduced the performance (i.e., the WER increased by about 2% relative) of
45 the tandem system compared to the phoneme-based ANNs trained on the same
46 language-independent data. Work by Lal and King (2013) compared the use
47 of AF-based ANNs trained on data from multiple languages (also including the
48 target-language) with AF-based ANNs trained on data from the target-language
49 data for tandem systems. It was observed that irrespective of whether the ANNs
50 are trained on data from multiple languages or on the target-language, the AF-
51 based ANNs performed better (relative improvement in WER of 1-9%) than
52 the phoneme-based ANNs. However, the AF-based (or phoneme-based) ANNs
53 trained on data from the target-language performed better than the AF-based
54 (or phoneme-based) ANNs trained on data from multiple languages (Lal and
55 King, 2013).
56
57
58

The difference in conclusion by Cetin et al. (2007b) and Lal and King (2013) could be because of the differences in the number of languages used to train the MLPs and their relationship with the target-language. Cetin et al. (2007b) used AF-based ANNs trained on English to generate tandem features for Mandarin ASR task; and English and Mandarin belong to different language families. Whereas Lal and King (2013) used ANNs trained on data from multiple languages that also included the target-language.

To summarize, most of the approaches to integrate AFs into an ASR system are based on the deterministic knowledge-based phoneme-to-AF relationship. The approaches summarized in Section 3.1 use the knowledge-based phoneme-to-AF relationship to define the deterministic lexical model parameters. The approaches described in Section 3.2 allow asynchronous evolution of various AFs using an AF-based pronunciation lexicon. However, the AF-based pronunciation lexicon is prepared using the knowledge-based phoneme-to-AF relationship.

In the next section, we present an ASR approach that integrates a model for lexical access using AFs into the HMM-based ASR framework. The approach adapts the knowledge-based phoneme-to-AF relationship using transcribed speech data and incorporates a probabilistic phoneme-to-AF relationship in the model parameters.

4. Proposed Approach

In probabilistic lexical model based ASR systems, a lexical unit is probabilistically related to all acoustic units (Section 2.4.2). For each lexical unit l^i , let $\mathbf{y}_i$ be the D -dimensional probability vector or the categorical variable that captures a probabilistic relationship between the lexical unit l^i and D acoustic units, i.e.,

$$\mathbf{y}_i = [y_i^1, \dots, y_i^d, \dots, y_i^D]^T \text{ where } y_i^d = P(a^d | l^i) \quad (13)$$

Therefore, the lexical model parameter set $\theta_l = \{\mathbf{y}_i\}_{i=1}^I$. In this paper we use the KL-HMM approach to estimate the parameters of the lexical model (Aradilla et al., 2007, 2008; Aradilla, 2008).

4.1. KL-HMM Approach

The KL-HMM approach for lexical model parameter estimation is summarized below:

- The approach assumes that the acoustic unit set $\mathcal{A}$ is known and a trained acoustic model is available. It has been shown that the acoustic units can be modeled using an ANN (Aradilla et al., 2007, 2008; Rasipuram and Magimai-Doss, 2015) or using GMMs (Rasipuram and Magimai-Doss, 2013a). In this paper, we use an ANN as the acoustic model.
- Given the acoustic model, the probabilities of acoustic units or the acoustic unit posterior probability vectors for the training data are estimated. At

time t , the acoustic unit posterior probability vector $\mathbf{z}_t$ is a D dimensional probability vector.

$$\begin{aligned}\mathbf{z}_t &= [z_t^1, \dots, z_t^d, \dots, z_t^D]^T \\ &= [P(a^1|\mathbf{x}_t), \dots, P(a^d|\mathbf{x}_t), \dots, P(a^D|\mathbf{x}_t)]^T\end{aligned}\quad (14)$$

- The acoustic unit probability vector sequences are used along with the pronunciation lexicon and word level transcriptions to train the parameters of the probabilistic lexical model. More precisely, the acoustic unit probability vector sequences are used as feature observations to train an HMM where the states represent the lexical units. Each state l^i is parameterized by a categorical variable $\mathbf{y}_i$ that captures a probabilistic relationship between a lexical unit l^i and D acoustic units.
- As both feature observations and state distributions are probability vectors, the local score at each HMM state can be computed as the KL-divergence between the feature observation $\mathbf{z}_t$ and the categorical variable $\mathbf{y}_i$. KL-divergence being an asymmetric measure, there are the following three possible ways to estimate the KL-divergence:

1. KL-divergence (S_{KL}): In this case, the state distribution $\mathbf{y}_i$ is the reference distribution.

$$S_{KL}(\mathbf{y}_i, \mathbf{z}_t) = \sum_{d=1}^D y_i^d \log \left(\frac{y_i^d}{z_t^d} \right) \quad (15)$$

2. Reverse KL-divergence (S_{RKL}): In this case, the acoustic unit probability vector $\mathbf{z}_t$ is the reference distribution.

$$S_{RKL}(\mathbf{y}_i, \mathbf{z}_t) = \sum_{d=1}^D z_t^d \log \left(\frac{z_t^d}{y_i^d} \right) \quad (16)$$

3. Symmetric KL-divergence (S_{SKL}): The local score S_{SKL} is the average of the local scores S_{KL} and S_{RKL} .

$$S_{SKL}(\mathbf{y}_i, \mathbf{z}_t) = \frac{1}{2} \cdot [S_{KL}(\mathbf{y}_i, \mathbf{z}_t) + S_{RKL}(\mathbf{z}_t, \mathbf{y}_i)] \quad (17)$$

- The parameters $\{\mathbf{y}_i\}_{i=1}^I$ are trained using the Viterbi training algorithm optimizing a function based on one of the KL-divergence based local scores.
- Decoding is performed using the standard Viterbi decoder where the log-likelihood based score is replaced with the KL-divergence based local score.

The details of the parameter estimation are elaborated in Appendix A. The details on the role of different local scores in estimating the lexical model parameters can be found in the work by Rasipuram and Magimai-Doss (2013b) whereas the role of local scores during decoding can be found in the work by Rasipuram and Magimai-Doss (2015).

4.2. Integrating AFs using KL-HMM

In the proposed approach, the relationship between lexical units and acoustic features is factored into two parts through the use of AFs as latent variables or acoustic units:

1. *The acoustic model* where the relationship between AFs and acoustic features is modeled.
2. *The lexical model* where a probabilistic relationship between AFs and lexical units is modeled.

The proposed approach exploits the advantage of probabilistic lexical modeling that the subword unit set used for defining the acoustic units need not be the same as the subword unit set used for defining the lexical units (Section 2.5). The lexical units can be based on context-independent or context-dependent subword units. The proposed approach for AF integration can be summarized in the following steps:

- The acoustic unit set consists of AFs such as manner and place of articulation. Therefore, the acoustic unit set can be seen as a superset of the individual AF sets, i.e.,

$$\mathcal{A} = \{\{\mathcal{A}_1\}, \dots, \{\mathcal{A}_F\}\} \quad (18)$$

where $\{\mathcal{A}_1\}, \dots, \{\mathcal{A}_F\}$ denote the individual AF sets and F the total number of AFs. For example, the set $\{\mathcal{A}_1\}$ may consist of all the classes specifying the manner of articulation such as vowel, stop, fricative, and so on; the set $\{\mathcal{A}_2\}$ may consist of all the classes specifying the place of articulation such as alveolar, back, dental, dorsal, front and so on. The total number of acoustic units

$$D = D_1 + \dots + D_F \quad (19)$$

where $D_1, \dots, D_F$ denote the cardinality of the individual AFs.

- Each AF is associated with an acoustic model, in our case, with an ANN.
- Given the AF-based acoustic models, posterior probabilities of AFs are estimated. The posterior probability estimates of various AFs are concatenated to produce a D dimensional acoustic unit probability vector $\mathbf{z}_t$, i.e.,

$$\mathbf{z}_t = [\mathbf{z}_{t,1}, \dots, \mathbf{z}_{t,F}]^T \text{ where} \quad (20)$$

$$\mathbf{z}_{t,1} = [z_{t,1}^1, \dots, z_{t,1}^{D_1}] \text{ similarly } \mathbf{z}_{t,F} = [z_{t,F}^1, \dots, z_{t,F}^{D_F}]$$

- The lexical model parameters $\theta_l = \{\{\mathbf{y}_{i,f}\}_{f=1}^F\}_{i=1}^I$ where $\mathbf{y}_{i,f}$ captures a probabilistic relationship between the lexical unit l^i and the AF classes

$\{\mathcal{A}_f\}$. That is,

$$\mathbf{y}_{i,1} = [y_{i,1}^1, \dots, y_{i,1}^{D_1}], \sum_{d=1}^{D_1} y_{i,1}^d = 1$$

$$\mathbf{y}_{i,F} = [y_{i,F}^1, \dots, y_{i,F}^{D_F}], \sum_{d=1}^{D_F} y_{i,F}^d = 1$$

- The parameters of the lexical model are trained using the KL-HMM approach. For a lexical unit l^i , the state distribution $\mathbf{y}_i$ can be seen as a stack of F categorical variables, i.e.,

$$\mathbf{y}_i = [\mathbf{y}_{i,1}, \dots, \mathbf{y}_{i,F}]^T \quad (21)$$

- The local score at each HMM state is the KL-divergence between the feature observation and the state distribution. If the local score is S_{RKL} , then the KL-divergence is computed as:

$$S_{\text{RKL}}(\mathbf{y}_i, \mathbf{z}_t) = \sum_{d=1}^{D_1} z_{t,1}^d \log \left(\frac{z_{t,1}^d}{y_{i,1}^d} \right) + \dots + \sum_{d=1}^{D_F} z_{t,F}^d \log \left(\frac{z_{t,F}^d}{y_{i,F}^d} \right) \quad (22)$$

Figure 3 illustrates the proposed AF-based ASR approach. The ANNs are trained to classify various AFs. The posterior probabilities of the individual AFs are stacked and used as feature observations to train an HMM. As mentioned in Section 2.2, each context-independent or context-dependent subword unit is modeled with three-HMM states.

4.3. Previous findings

The proposed approach for AF integration was studied for phoneme recognition using the TIMIT corpus (Rasipuram and Magimai-Doss, 2011b,a,c). In our previous work the notion of probabilistic lexical modeling was not introduced and the studies were presented from the perspective of acoustic modeling. However, as shown in our recent work, KL-HMM is a lexical modeling approach (Rasipuram and Magimai-Doss, 2013a, 2015). Given the formulation of Section 4.2, the findings from the previous studies are refined and re-summarized below:

- It was demonstrated that using AF-based acoustic models in the KL-HMM approach results in better system than using the same acoustic models in the hybrid HMM/ANN approach (Rasipuram and Magimai-Doss, 2011b). The study illustrated that it is beneficial if the lexical model is probabilistic.
- Performance of the KL-HMM system using both phoneme-based and AF-based acoustic models was better than the KL-HMM system using only the phoneme-based acoustic model. The study indicates that the KL-HMM approach has the potential to reduce the error rates by incorporating articulatory knowledge into an ASR system.

Figure 3: Illustration of the KL-HMM approach for AF integration.

- The phoneme recognition performance of a system can be improved by improving the AF acoustic models. It was observed that the AF classification accuracy can be improved by modeling the inter-feature dependencies using multistage MLP classifiers and/or multitask learning. In doing so, the performance gap between systems using phonemes as acoustic units and AFs as acoustic units was greatly reduced.

4.4. Other Contributions of the Paper

The main contribution of the paper is the formulation of AF-based speech recognition in the framework of probabilistic lexical modeling. In addition to that and the previous findings summarized in Section 4.3, the contributions of the paper are as follows:

- The proposed approach for AF integration is evaluated on a continuous speech recognition task (see Section 5 for the experimental setup). In the evaluation, acoustic models are ANNs estimating various multivalued AFs. The lexical units are based on context-dependent phonemes and the lexical model is probabilistic. The proposed approach is compared with the tandem approach for AF integration (Section 6.1). The tandem

systems also use the same ANNs as the KL-HMM systems, but as feature extractors.

- Exploiting the resource optimization advantage of probabilistic lexical modeling (see Section 2.5), we study the case where AF-based acoustic models are trained on domain-independent resources and the lexical model parameters are trained on domain-dependent resources (see Section 6.2). Furthermore, the use of multistage MLP classifiers is studied for a continuous speech recognition task (Section 6.3).
- In the framework of probabilistic lexical modeling it is possible to build grapheme-based ASR systems where the lexical units are based on graphemes and the acoustic units are based on phonemes (Magimai-Doss et al., 2011; Rasipuram and Magimai-Doss, 2013a, 2015). These grapheme-based ASR systems, where the lexical model parameters capture a probabilistic graphemes-to-phoneme relationship, performed better than the grapheme-based ASR approaches where the lexical model is deterministic. Motivated by this, in this paper we hypothesize that it is possible to build grapheme-based ASR systems where the lexical units are based on graphemes and the acoustic units are AFs. In this case, the lexical model parameters capture a probabilistic grapheme-to-AF relationship. The resulting grapheme-based ASR approach, in addition to exploiting the advantages of AFs, can also address resource constrained speech recognition with limited or no pronunciation lexicon.
- The lexical model parameters capture a probabilistic relationship between phonemes and AFs. We analyze the parameters of the lexical model to understand if the captured phoneme-to-AF relationship associates well with the knowledge-based phoneme-to-AF relationship (see Section 7).

5. Experimental Studies

In this paper, we evaluate the proposed approach for AF integration on a speaker-independent continuous speech recognition task using the DARPA Resource Management (RM) corpus.

5.1. Database

The RM corpus consists of read queries on the status of Naval resources (Price et al., 1988). The task is artificial in aspects such as speech type, range of vocabulary and grammatical constraint. The training set consists of 2880 utterances and the development set 1110 utterances. The training and development sets together consist of 3,990 utterances spoken by 109 speakers corresponding to approximately 3.8 hours of speech data.

There are four test sets provided by DARPA, namely, feb89, oct89, feb91, and sep92. Each test set contains 300 utterances spoken by 10 speakers. The test set used in this work is obtained by combining the four test sets and thus contains

1,200 utterances amounting to 1.1 hours in total. The test set is completely covered by a word pair grammar included in the task specification.

The phoneme-lexicon consists of 991 words. In this paper, we followed the RM setup used by Dines and Magimai-Doss (2007) for the pronunciation lexicon to have fairer comparison with the previous work. The UNISYN² lexicon with general American accent was used. There are 45 context-independent phonemes including silence. For ASR experiments, the 45 phonemes were converted to a set of 42 phonemes by merging some allophones (/em/, /en/ and /el/) to their broader phoneme class (/m/, /n/ and /l/). About 35 words in the phoneme lexicon have more than one pronunciation.

The grapheme-lexicon for the RM task is transcribed using 79 graphemes. The first grapheme and the last grapheme of a word are treated as separate graphemes. Therefore, the grapheme set consists of 26 English graphemes ({[A],[B],...[Z]}), 26 English graphemes occurring at the begin of word ({[b_A],[b_B],...[b_Z]}), 26 English graphemes occurring at the end of word ({[e_A],[e_B],...[e_Z]}) and silence.

5.2. MLPs

KL-HMM systems use ANNs, more specifically, multilayer perceptrons (MLPs) as the acoustic models whereas tandem systems use the same MLPs as feature extractors. The input to all the MLPs is 39-dimensional perceptual linear prediction (PLP) features with a nine frame temporal context (i.e., four frame preceding and four frame following context). We use three-layer (one input, one hidden and one output layer) MLPs that are trained with the frame level cross entropy error criteria using the Quicknet software³. The number of hidden units of the MLPs are selected based on the optimal frame accuracy on the development set.

The target labels for the MLPs with phonemes as output units were obtained from the HMM/GMM system. The target labels for the MLPs with AFs as output units are obtained from the phoneme-to-AF map given in Table B.1. The AFs consist of manner, place, height, and vowel. The phoneme-to-AF map is adapted (to distinguish all phonemes of the RM task) from the mapping defined by Hosom (2009). The place class is expanded by adding features like mid-front and mid-back. The height class is expanded by adding features like mid, mid-low, mid-high. Also, the vowel AF is added.

5.2.1. Domain-dependent MLPs

We use the following MLPs trained on the RM corpus:-

1. *MLP-RM-PH*: An *off-the-shelf* MLP (Dines and Magimai-Doss, 2007) trained on the RM corpus to classify 45 context-independent phonemes.
2. *MLP-RM-AF*: Set of MLPs trained on the RM corpus to classify AFs.

²<http://www.cstr.ed.ac.uk/projects/unisyn/>

³<http://www.icsi.berkeley.edu/Speech/qn.html>

3. *MLP-RM-PH+AF*: The phoneme and the AF MLPs together are referred as *MLP-RM-PH+AF*.

5.2.2. Domain-independent MLPs

Exploiting the resource optimization advantage of probabilistic lexical modeling (Section 2.5) we study the case where the acoustic model is trained on domain-independent data whereas the lexical model is trained on target-domain data. This also allows us to study the data-invariance aspect of AFs. Similarly, a tandem system can use the MLP trained on domain-independent data for feature extraction. In this paper, we use the Wall Street Journal (WSJ) corpus (Paul and Baker, 1992; Woodland et al., 1994) as domain-independent data to train the MLPs whereas the RM corpus is used as domain-dependent data for which we are interested to build an ASR system. The WSJ corpus consists of two parts - WSJ0 with 14 hours of speech (7,193 utterances from 84 speakers) and WSJ1 with 66 hours of speech (29322 utterances from 200 speakers). In this paper, we use only the WSJ1 corpus as the domain-independent data. Among 1000 words present in the RM task, WSJ corpus consists of only 568 words. The phoneme-lexicon for the WSJ corpus was also obtained from the UNISYN lexicon. Therefore, the phoneme sets and the AF sets for the RM and WSJ corpora are identical. In this paper, we use the following MLPs trained on the WSJ corpus:-

1. *MLP-WSJ-PH*: An *off-the-shelf* MLP (Aradilla et al., 2008) trained on the WSJ corpus to classify 45 context-independent phonemes.
2. *MLP-WSJ-AF*: Set of MLPs trained on the WSJ corpus to classify AFs.
3. *MLP-WSJ-PH+AF*: The phoneme and articulatory MLPs together are referred as *MLP-WSJ-PH+AF*.

5.2.3. Multistage MLPs

In this paper, we use the multistage MLP classifier illustrated in Figure 4. In the first stage, a set of parallel MLPs are used to estimate posteriors of the AFs. Each MLP receives PLP features as input and is trained to classify a specific AF. In the second stage, to model the temporal contextual information and inter-feature dependencies of AFs, a new set of MLPs are trained using articulatory posteriors estimated by the first stage of MLPs (along with other AFs) and a longer temporal context (eight frame preceding and eight frame following context) as input. In this paper, we use the following multistage MLPs:-

1. *MULTI-WSJRM-PH*: The phoneme-based multistage MLP, where the first stage MLP is trained on the WSJ corpus to classify context-independent phonemes and the second stage MLP is trained on the RM corpus to classify context-independent phonemes.
2. *MULTI-WSJRM-AF*: Set of AF-based multistage MLPs as shown in Figure 4, where the first set of AF-based MLPs are trained on the WSJ corpus and the second set of AF-based MLPs are trained on the RM corpus.
3. *MULTI-WSJRM-PH+AF*: The phoneme and articulatory multistage MLPs together are referred as *MULTI-WSJRM-PH+AF*.

Figure 4: Multistage AF-based MLP classifiers.

The input, output and hidden layer sizes of all the MLPs used in the paper along with the total number of parameters and frame accuracies are given in Table B.2 of Appendix.

5.3. Systems

We compare the KL-HMM, tandem and HMM/GMM systems trained on both the training and development sets (3,990 utterances) of the RM corpus. All the systems model crossword context-dependent subword units (either phonemes or graphemes) and each subword unit is modeled as a 3-state HMM. Table 1 summarizes the different systems and their capabilities.

KL-HMM systems: The acoustic units in the KL-HMM system can be context-independent phonemes, or AFs or both context-independent phonemes and AFs. KL-HMM systems use an MLP as the acoustic model. The lexical units are based on context-dependent phonemes and the lexical model is probabilistic. The lexical model is trained using the local score (S_{SKL} or S_{RKL} or S_{KL}) that results in minimum KL-divergence on the training data compared to other local scores. It was observed that for phoneme-based KL-HMM systems, the local score that resulted in minimum KL-divergence was S_{SKL} whereas for grapheme-based KL-HMM systems it was either S_{SKL} or S_{RKL} .

Tandem systems: The tandem systems use the MLPs as feature extractors. The MLPs used for feature extraction are the same MLPs that are used as the acoustic model in the KL-HMM systems. The output of the MLPs is Gaussianized with log transformation followed by KLT. The dimensionality of the features is reduced by retaining the feature components that contribute to 99% of the variance. The resulting features are used to train an HMM/GMM system where the acoustic units are the clustered context-dependent subword states and the lexical units are based on the context-dependent subword units. The lexical and acoustic units are deterministically related. Each acoustic unit is modeled with an eight component Gaussian mixture model as it resulted in an optimal ASR performance.

HMM/GMM systems: The 39-dimensional PLP features used to train the MLP are also used to train the HMM/GMM systems where the acoustic units are clustered context-dependent subword states and the lexical units are based on

context-dependent subword units. The lexical and acoustic units are deterministically related. Each acoustic unit is modeled with an eight mixture Gaussian. The state tying resulted in 1611 tied states for the phoneme-based system and 1536 tied states for the grapheme-based system.

The details of various systems, such as the number of lexical units, acoustic units, dimensionality of tandem features, tied states in tandem systems, total number of parameters in different systems are given in Table B.3.

Table 1: Overview of different systems. In the table *ph*, *af* and *af+ph* denote that acoustic units are phonemes, AFs, and both AFs and phonemes, respectively. cCD denotes clustered context-dependent subword states. In the tandem approach, the ANNs trained to classify the acoustic units (*ph*, *af* or *af+ph*) are used to extract features for the HMM/GMM system. This is indicated through (ANN) notation. *Det* and *Prob* denote that the lexical model is deterministic and probabilistic, respectively.

System	Acoustic units	Acoustic model	Lexical units	Lexical model
KL-HMM	{ph}, or {af} or {af+ph}	ANN	CD	Prob
Tandem	cCD	(ANN) + GMM	CD	Det
HMM/GMM	cCD	GMM	CD	Det

6. Results

The performance of the phoneme-based and grapheme-based KL-HMM and tandem systems using the various MLPs in terms of word accuracy on the test set of the RM corpus is given in Table 2. Performance on the test set of the RM corpus for the standard crossword context-dependent phoneme-based HMM/GMM system is 95.9% word accuracy and the crossword context-dependent grapheme-based HMM/GMM system is 94.8% word accuracy. The performances of the baseline phoneme-based HMM/GMM systems of this paper and the phoneme-based HMM/GMM systems in the literature reported by Hain and Woodland (1999) and Povey et al. (2011b) on the RM corpus are the same despite the difference in the phoneme-lexicon used. In this work, we used a phoneme-lexicon based on the UNISYN lexicon whereas Hain and Woodland (1999) used a phoneme-lexicon based on the CMUDict.

In the remainder of the section, the results of the phoneme and grapheme-based KL-HMM and tandem systems are analysed according to the various MLPs used (as discussed in Section 5.2).

6.1. Domain-dependent MLPs

For both phoneme and grapheme subword units, the results indicate that:

Table 2: Performance in terms of word accuracy on the test set of the RM corpus for the crossword context-dependent phoneme-based and the grapheme-based KL-HMM and tandem systems. The MLPs *MLP-RM-PH*, *MLP-RM-AF* and *MLP-RM-PH+AF* are trained on the RM corpus. The MLPs *MLP-WSJ-PH*, *MLP-WSJ-AF* and *MLP-WSJ-PH+AF* are trained on the domain-independent WSJ corpus. The multistage MLPs *MULTI-WSJRM-PH*, *MULTI-WSJRM-AF* and *MULTI-WSJRM-PH+AF* are trained on both RM and WSJ corpora. Performance on the test set of the RM corpus for the standard cross-word context-dependent phoneme and grapheme subword based HMM/GMM systems is 95.9% and 94.8% word accuracy, respectively.

MLPs		Phoneme-based		Grapheme-based	
		KL-HMM	Tandem	KL-HMM	Tandem
domain-dependent	<i>MLP-RM-PH</i>	95.6	95.7	95.1	94.6
	<i>MLP-RM-AF</i>	94.9	95.3	94.4	94.1
	<i>MLP-RM-PH+AF</i>	96.2	95.3	95.8	94.3
domain-independent	<i>MLP-WSJ-PH</i>	95.9	95.8	95.5	94.5
	<i>MLP-WSJ-AF</i>	95.5	95.4	94.9	94.8
	<i>MLP-WSJ-PH+AF</i>	96.4	94.9	96.0	94.5
multistage	<i>MULTI-WSJRM-PH</i>	96.1	96.0	95.8	95.6
	<i>MULTI-WSJRM-AF</i>	95.4	95.2	95.8	95.2
	<i>MULTI-WSJRM-PH+AF</i>	96.6	95.1	96.4	95.0

- The KL-HMM systems using the phoneme-based MLP (*MLP-RM-PH*) significantly ⁴ outperform the KL-HMM systems using the AF-based MLPs (*MLP-RM-AF*). The KL-HMM systems using both phoneme-based and AF-based MLPs (*MLP-RM-PH+AF*) significantly outperform the KL-HMM systems using either phoneme-based or AF-based MLPs.
- The difference in performance between the KL-HMM and tandem systems is significant when both phoneme-based and AF-based MLPs (*MLP-RM-PH+AF*) are used. However, the difference in performance between the KL-HMM and tandem systems employing either the phoneme-based (*MLP-RM-PH*) or the AF-based (*MLP-RM-AF*) MLPs is not significant.
- The difference in performance between the KL-HMM and tandem systems using the phoneme-based MLP (*MLP-RM-PH*), and the baseline HMM/GMM system is not statistically significant.

In this section, the AF-based MLPs are trained on a limited amount of domain-dependent data. In the next section, we verify that AF estimation using MLPs is data-invariant and can benefit from a larger domain-independent data set.

⁴In the paper, “significant difference” implies that the difference in performance between the compared systems is statistically significant with confidence greater than 95.0%. Statistical significance tests for the systems are performed using the approach proposed by Bisani and Ney (2004).

6.2. Domain-independent MLPs

In this case, the KL-HMM and tandem systems use the MLPs trained on the domain-independent WSJ corpus. The results show that:

- The phoneme-based KL-HMM system using the AF-based MLPs (*MLP-WSJ-AF*) trained on a large set of domain-independent data (the WSJ corpus) significantly outperforms the phoneme-based KL-HMM system using the AF-based MLPs (*MLP-RM-AF*) trained on a small set of domain-dependent data (the RM corpus). For all the other phoneme-based KL-HMM systems and all the grapheme-based KL-HMM systems, there is no significant difference in performance between the systems using domain-dependent MLPs and the systems using domain-independent MLPs.
- The grapheme-based tandem system using domain-dependent AF-based MLPs (*MLP-WSJ-AF*) performs significantly better than the grapheme-based tandem system using domain-independent AF-based MLPs (*MLP-RM-AF*). However, for all the other grapheme-based and all the phoneme-based tandem systems, there is no significant difference in performance between the systems using domain-dependent MLPs and the systems using domain-independent MLPs.
- Unlike the domain-dependent MLPs case, the difference in performance between the phoneme-based KL-HMM systems using the phoneme-based MLP (*MLP-WSJ-PH*) and the AF-based MLPs (*MLP-WSJ-AF*) is not statistically significant.
- Similar to the domain-dependent MLPs case, the phoneme-based and grapheme-based KL-HMM systems using both the phoneme-based and AF-based MLPs (*MLP-WSJ-PH+AF*) significantly outperform all other KL-HMM systems (employing other domain-independent MLPs) and all the tandem systems.

6.3. Multistage MLPs

In this case, the KL-HMM and tandem systems use the multistage MLPs trained on both WSJ and RM corpora. The results show that:

- The grapheme-based KL-HMM systems using either the AF-based multistage MLPs or both AF and phoneme-based multistage MLPs (i.e., *MULTI-WSJRM-AF* and *MULTI-WSJRM-PH+AF*) significantly outperform the grapheme-based KL-HMM systems using the respective single-stage MLPs trained on the WSJ corpus (*MLP-WSJ-AF* and *MLP-WSJ-PH+AF*). However, for all the phoneme-based KL-HMM systems and all other grapheme-based KL-HMM systems, there is no significant difference in performance between the systems using the multistage MLPs and the systems using domain-independent MLPs.

- Unlike the domain-dependent and domain-independent MLPs case, the difference in performance between the grapheme-based KL-HMM system using the phoneme-based multistage MLP (*MULTI-WSJRM-PH*) and the AF-based multistage MLPs (*MULTI-WSJRM-AF*) is not statistically significant.
- The performance gap between grapheme-based and phoneme-based KL-HMM systems is greatly reduced when multistage MLPs are used.

6.4. Summary of the Experimental Results

To summarize, the following conclusions can be drawn from the experimental study:

1. The proposed approach for AF integration resulted in an ASR system that performs similar to the tandem approach for AF integration. Though both approaches perform similarly, there are two main advantages of the proposed approach.
Firstly, in the proposed approach, the articulatory representations are kept intact in the model parameters in the form of probabilistic phoneme-to-AF or grapheme-to-AF relationship learned from the transcribed speech data. Furthermore, as we will see in the next section, the approach also adapts the knowledge-based phoneme-to-AF and grapheme-to-AF relationship on the transcribed speech data, and allows different AFs to evolve asynchronously. However, the tandem approach tends to lose the two primary benefits of articulatory representation, i.e., finer granularity and asynchronous evolution.
Secondly, the KL-HMM system achieves similar performance as the tandem system but uses fewer parameters than the tandem approach (about 30%-40% relative for the single-stage MLPs trained on the RM corpus, and about 10% relative for the single-stage MLPs trained on the WSJ corpus and multistage MLPs as observed in Table B.3).
2. The performance of the grapheme- or phoneme-based KL-HMM systems using both AFs and phonemes as acoustic units is always better (about 10 to 12% relative reduction in WER) than the KL-HMM system using either of them as acoustic units. We speculate the following reasons:
 - (a) When both AFs and phonemes are used as acoustic units, not only the acoustic model is improved but also the lexical model, as both probabilistic phoneme-to-phoneme and phoneme-to-AF (or grapheme-to-phoneme and grapheme-to-AF) relationships are modeled.
 - (b) The AF-based and phoneme-based MLPs are trained independently. However, the probabilistic phoneme-to-phoneme and phoneme-to-AF (or grapheme-to-phoneme and grapheme-to-AF) relationships are learned together during lexical model training. Therefore, the approach can learn the inter-feature dependencies among various AFs and phonemes or graphemes.
 - (c) The information captured by the lexical model with phonemes as acoustic units and with AFs as acoustic units is complementary.

The proposed approach significantly outperforms the tandem approach if both AFs and phonemes are used as acoustic units. The performance of the grapheme- or phoneme-based tandem systems using both AFs and phonemes as acoustic units is always worse than the tandem system using either of them as acoustic units. This could be because the statistical characteristics of phoneme and articulatory feature probabilities are different. It may be necessary to model phonemes and AFs separately such as in the factored observation model (Cetin et al., 2007a).

3. The results indicate that AFs estimation using MLPs is data and domain invariant. Further, the use of AF-based MLPs trained on a larger domain-independent data set helps both phoneme-based and grapheme-based KL-HMM systems.
4. In our previous phoneme recognition studies it was observed that the multistage articulatory MLPs improved the phoneme recognition accuracy (Rasipuram and Magimai-Doss, 2011b). In this paper, multistage AF-based MLPs did not improve the performance of the phoneme-based KL-HMM systems but improved the performance of the grapheme-based KL-HMM systems. We conjecture the following two reasons for the observed trends.

Firstly, the multistage AF-based MLP was motivated from the work by Pinto et al. (2011). In that work it was shown that the second MLP in the multistage MLP classifier learns the phonetic temporal patterns (i.e., the phonetic confusions at the output of the first MLP) and the phonotactics of the language observed in the training data. In our case, the second set of MLPs in the multistage AF-based MLP classifier could model phonotactic constraints at the articulatory feature level.

 - When the lexical units are based on context-dependent phonemes, the lexical model incorporates phonotactic constraints. The results indicate that it may be redundant to model the phonotactic constraints twice, once at the acoustic model level and again at the lexical model level.
 - In the case of context-dependent grapheme-based KL-HMM systems, the lexical model is modeling graphemic constraints and the acoustic model is modeling phonotactic constraints, which could be complementary to each other especially given the fact that the grapheme-to-phoneme relationship in English is irregular.

Secondly, though the frame accuracy and phoneme recognition accuracy are considered as important factors for word recognition, they may not be the only indicators of word level performance (Greenberg et al., 2000). The relationship between frame accuracy and word accuracy depends on more than one factor: the pronunciation lexicon, the acoustic model, the lexical model and the language model components of an ASR system. In other words, lexical constraints and syntactic constraints incorporated while decoding can handle the shortcomings of the acoustic model or may render some of the gains obtained through the acoustic model redundant.

7. Analysis

In the proposed approach, with phonemes or graphemes as lexical units and AFs as acoustic units, the parameters of the lexical model capture a probabilistic phoneme-to-AF or grapheme-to-AF relationships. In this section, we analyze the parameters of the lexical model at subword level (Section 7.1) and word level (Section 7.2) to understand the following:

- Is the phoneme-to-AF or grapheme-to-AF relationship captured by the lexical model parameters close to the knowledge-based phoneme-to-AF or grapheme-to-AF relationship?
- Does the model allow different AFs to evolve asynchronously?

The target labels for the MLPs with AFs as output units are obtained from the knowledge-based phoneme-to-AF map. Hence, during training, the MLPs do not account for the AFs changing asynchronously. However, as shown by King and Taylor (2000), it is typical that at the output of the MLP, different AFs change at different times (especially at the phoneme boundaries) and exhibit asynchronous behaviour. Since the KL-HMM system models the output of a set of AF-based MLPs, we hypothesize that the lexical model parameters can capture asynchronous AF behaviour especially at the HMM states modeling subword unit boundaries (i.e., the first and third HMM states of a subword unit).

7.1. Subword Level Analysis

Table 3 shows the manner and place of articulation for context-independent phonemes in three HMM states captured by the lexical model parameters. The analysis is presented only on manner and place of articulation for the sake of simplicity. However, similar trends are observed even for height of articulation. The denoted AF values correspond to the dimension with maximum probability captured by the lexical model parameters of context-independent subword units. The first and second parts of the table presents examples of context-independent phonemes where the manner and place of articulation between three HMM states are synchronous and asynchronous, respectively. It can be observed that the phoneme-to-AF relationship of Table 3 relates well with the knowledge-based relationship given in Table B.1.

Table 3 indicates that for context-independent phonemes that are diphthongs such as /aw/, /ay/, /ow/, and /oy/ and for phonemes that are vowels such as /ah/ and /uh/ the captured place of articulation is the same, i.e., “vowel”. For phonemes that are voiced-stops, i.e., /b/, and /g/, the captured place of articulation is the same across the three HMM states whereas the captured manner of articulation changed between the HMM states. More specifically, the initial states of /b/ and /g/ captured a “voiced-stop” and the third state captured a “vowel”. For phonemes /ch/ and /jh/, the captured manner of articulation changed at the second HMM state whereas the captured place of articulation changed at the first HMM state.

Table 3: The phoneme-to-manner of articulation and phoneme-to-place of articulation relationship captured by the lexical model parameters of context-independent phonemes.

Synchronous Examples				
Phoneme	AF	state1	state2	state3
aa	Manner	vow		
	Place	bck		
d	Manner	vst		
	Place	alv		
l	Manner	app		
	Place	lat		
m	Manner	nas		
	Place	lab		
Asynchronous Examples				
Phoneme	AF	state1	state2	state3
aw	Manner	vow		
	Place	midf	bck	midb
ay	Manner	vow		
	Place	bck	midf	
ow	Manner	vow		
	Place	bck		midb
oy	Manner	vow		
	Place	bck	lat	midf
ah	Manner	vow		
	Place	bck	mid	
uh	Manner	vow		
	Place	bck	midb	
b	Manner	sil	vst	vow
	Place	lab		
g	Manner	vst		vow
	Place	dor		
ch	Manner	stp		frc
	Place	alv	fnt	
jh	Manner	vst		vow
	Place	alv	fnt	
er	Manner	app		vow
	Place	ret		
ng	Manner	nas		
	Place	dor		alv

We have computed the percentage of context-dependent phonemes where the lexical model parameters exhibited asynchrony between manner and place of articulation at the HMM state level. A context-dependent phoneme model is said to be synchronous if manner and place of articulation change at the same state transition or remain the same across three HMM states. A context-dependent phoneme model is said to be asynchronous if manner and place of articulation change at different HMM states. For example, in Table 4, the context-dependent phonemes /aa-n+t/ and /ah-z+sh/ exhibit synchronous changes in manner and place of articulation whereas the phonemes /w-ey+r/ and /sil-jh+ae/ exhibit asynchronous changes at the HMM state level.

In Table 5, the first column indicates the set of MLPs used to train the KL-HMM system, and the second and the third columns indicate the percentage of

Table 4: Examples of the context-dependent phonemes that exhibit synchronous and asynchronous manner and place of articulation at the HMM state level.

Synchronous Examples					Asynchronous Examples				
Phoneme	AF	st1	st2	st3	Phoneme	AF	st1	st2	st3
aa-n+t	Manner	nas			w-ey+r	Manner	vow		
	Place	alv				Place	bck	midf	ret
ah-z+sh	Manner	vow	frc		sil-jh+ae	Manner	sil	vst	vow
	Place	midf	frt			Place	sil	fnt	

context-dependent phonemes where the changes in manner and place of articulations at the HMM state level are synchronous and asynchronous, respectively. It is important to note that the classification of context-dependent phoneme models in terms of synchronous and asynchronous does not take into account the errors in the phoneme-to-AF map captured by the lexical model parameters. For example, in Table 7, the place and height of articulation for the grapheme model [W] are asynchronous. However, the captured place of articulation in the first state of [W] as “lateral” could be considered as an error. Therefore, the percentage of context-dependent phoneme models exhibiting synchronous or asynchronous behaviour are only an indicative of the asynchronous nature of the context-dependent phoneme models.

Table 5 indicates that asynchronous articulatory movements among manner and place of articulation are relatively lower when multistage AF-based MLP classifiers are used. We argue the following two reasons for this. Firstly, as discussed in Section 6.4, the second stage of AF-based MLPs in the multistage MLP classifiers model the phonotactics of the language, therefore various AFs may be more synchronous. Secondly, the frame accuracies of the multistage AF-based MLP classifiers are better than the frame accuracies of other MLPs. As a result the number of context-dependent phonemes exhibiting asynchronous changes because of the errors in the captured relationship could be relatively less.

Table 5: For various KL-HMM systems, the percentage of context-dependent phonemes where the changes in manner and place of articulations between the three HMM states are synchronous and asynchronous.

MLP	Synchronous	Asynchronous
<i>MLP-RM-AF</i>	59.49%	40.50%
<i>MLP-WSJ-AF</i>	61.56%	38.40%
<i>MULTI-WSJRM-AF</i>	67.50%	32.49%

The analysis of the parameters indicated that the model is able to capture asynchronous AF configurations at the subword unit level. In the next section we will see that asynchronous articulatory configurations are more meaningful at the word level as the context-dependent subword models also capture information of the neighbouring phonemes.

Table 6: The phoneme-to-AF relationship captured by the lexical model parameters of context-dependent phonemes for the word “BELOW”. Pronunciation of the word BELOW in the phoneme lexicon is /b/ /ih/ /l/ /ow/. AF values for manner, place and height of articulation are given.

Phoneme	b			ih			l			ow		
State	st1	st2	st3	st1	st2	st3	st1	st2	st3	st1	st2	st3
Manner	sil	vst	vow				app			vow		
Place	sil	lab		midf		lat				bck		midb
Height	max		high		vhi				mid		vhi	

Table 7: The grapheme-to-AF relationship captured by the lexical model parameters of context-dependent graphemes for the word “BELOW”. Pronunciation of the word BELOW in the grapheme lexicon is [B] [E] [L] [O] [W]. AF values for manner, place and height of articulation are given.

Grapheme	[B]			[E]			[L]			[O]			[W]		
State	st1	st2	st3	st1	st2	st3	st1	st2	st3	st1	st2	st3	st1	st2	st3
Manner	vst		vow		app				vow						
Place	sil	lab		midf		lat			bck		<i>lat</i>	midb	sil		
Height	max		high		vhi		mid		<i>midl</i>		vhi		sil		

7.2. Word Level Analysis

The phoneme-to-AF and grapheme-to-AF relationships captured by the lexical model parameters of phoneme-based and grapheme-based KL-HMM systems for the word “BELOW” are given in Tables 6 and 7, respectively. The tables indicate the manner, place and height of articulation.

It can be observed that the phoneme-to-AF relationship of Table 6 relates well with the knowledge-based relationship given in Table B.1. Table 7 shows that even if subword units are graphemes, articulatory patterns similar to the system using phoneme subword units are captured. The two differences between grapheme-to-AF and phoneme-to-AF are indicated in red italic font in Table 7. The number of subword units in the pronunciation of the word “BELOW” are five in the case of graphemes and four in the case of phonemes. It can be observed from the table that this irregularity in the grapheme pronunciation has been accounted for, as the sequence of graphemes [O] and [W] together capture the information of phoneme /ow/.

Furthermore, the tables also indicate that various AFs evolve asynchronously. For example, in Table 6, the captured manner of articulation in the second state of /b/ is “voiced-stop” and in the third state of /b/ is “vowel”, whereas the place of articulation in both second and third states of /b/ is “labial”.

It was observed by Magimai-Doss et al. (2011) and Rasipuram (2014)[Chapter 4] that the lexical model parameters of the context-dependent graphemes tend to model the transition information to the next context-dependent grapheme in the sequence. Similar observations can also be made from Tables 6 and 7, but at the finer articulatory feature level. For example, in Table 6, the captured manner of articulation in the third state of /b/ is “vowel” which corresponds to the next phoneme in the sequence, i.e., /ih/. Similarly in Table 7, the captured manner, place and height of articulation in the third state

of grapheme [B] correspond to the next grapheme in the sequence, i.e., [E]. The analysis shows that the lexical model parameters of the context-dependent phonemes and graphemes are capable of capturing some information about preceding and following articulatory configurations.

8. Discussion and Conclusion

In this paper, we proposed an approach to integrate articulatory feature representations into HMM-based ASR in the framework of probabilistic lexical modeling. The proposed approach involves two stages: acoustic model and lexical model. The acoustic model is a posterior probability estimator that models the relationship between acoustic feature observations and AFs. The lexical model, models a probabilistic relationship between the lexical units and the AFs. The approach has the following potential advantages:

- Lexical access: As opposed to knowledge-based approaches, the parameters of the lexical model in the proposed approach are learned using transcribed speech data by training an HMM whose states represent lexical units and the parameters of the state capture a probabilistic relationship between a lexical unit and AFs. Consequently, the approach integrates a model for lexical access using AFs into the HMM-based ASR framework.
- Asynchrony of AFs: As observed in Section 7, the model also allows different AFs to evolve asynchronously. Thus, overcoming some of the limitations of knowledge-based approaches.
- Combination of various AFs: A challenge often faced in using articulatory features for ASR is the combination of evidences from different AFs. In that regard, the proposed approach can be seen as a multi-channel approach where each AF serves as a separate channel and various AFs are combined at the local score computation level. Also, as seen in this paper, the multi-channel approach can be trivially extended to combine other relevant information such as the phoneme information.

Our investigations on a continuous speech recognition task have shown that the proposed approach effectively integrates AFs into the HMM-based ASR framework; improves ASR performance if combined with phoneme-based acoustic models; exploits domain-independent resources; and offers flexibility to use either phonemes or graphemes as subword units.

The probabilistic grapheme-to-AF relationship captured in the lexical model parameters of the KL-HMM system with acoustic units as AFs and lexical units based on context-dependent graphemes can be exploited to generate an AF-based pronunciation lexica using the acoustic data-driven grapheme-to-phoneme conversion approach proposed by Rasipuram and Magimai-Doss (2012). The AF-based pronunciation lexica can be used in DBN-based approaches for AF integration that require such lexica (Livescu and Glass, 2004; Livescu et al., 2008) or in AF-based text-to-speech synthesis systems.

In this paper, we focussed mainly on the integration of AFs into an ASR system. More precisely, we focussed on the lexical model aspect of the proposed approach. The three-layer or multistage MLPs classifying context-independent AFs (or phonemes) were used as acoustic models. The approach can be potentially improved by improving the acoustic model along the following directions:

1. Context-dependent AFs: The output of MLPs or the acoustic units could be context-dependent AFs that take into account the neighbouring articulatory context.
2. Deep architectures for AF estimation: More recently, ANNs with deep architectures have gained lot of attention (Dahl et al., 2012; Hinton et al., 2012). In similar a vein, the articulatory feature model can be based on deep ANN architectures (Siniscalchi et al., 2012).

In this paper, we have shown that the AF-based acoustic models can be trained on domain-independent data whereas the lexical model can be trained on domain-dependent data. In our recent work we found that in the framework of probabilistic lexical modeling, the acoustic model can be trained on language-independent resources and the lexical model on a relatively small amount of language-dependent data (Rasipuram and Magimai-Doss, 2015). AFs are considered to be more language-independent and effective for cross-linguistic adaptation (Lal and King, 2013; Siniscalchi et al., 2012). Therefore, we hypothesize that the use of articulatory feature based language-independent acoustic model in the proposed approach can offer potential advantages in building ASR systems for under-resourced and minority languages⁵. Our future work will focus on extending the proposed approach along this direction.

Acknowledgment

This work was partly supported by the Swiss NSF through the grants “Flexible Grapheme-Based Automatic Speech Recognition (FlexASR)” and partly by the Commission for Technology and Innovation (CTI) on “Automatic scoring and adaptive pedagogy for oral language learning (ScoreL2)”.

References

- Aradilla, G., 2008. Acoustic Models for Posterior Features in Speech Recognition. Ph.D. thesis. EPFL. Switzerland.
- Aradilla, G., Boulard, H., Magimai-Doss, M., 2008. Using KL-Based Acoustic Models in a Large Vocabulary Recognition Task , in: Proc. of Interspeech, pp. 928–931.

⁵A language that lacks one or more resources required to build an ASR system is referred to as under-resourced language (Besacier et al., 2014). Minority languages are languages spoken by a minority of population. A minority language need not be under-resourced and an under-resourced language may or may not be a minority language.

- Aradilla, G., Vepa, J., Bourlard, H., 2007. An Acoustic Model Based on Kullback-Leibler Divergence for Posterior Features, in: Proc. of ICASSP, pp. 657–660.
- Bahl, L., Brown, P., de Souza, P., Picheny, M., 1988. Acoustic Markov models used in the Tangora speech recognition system, in: Proc. of ICASSP, pp. 497–500 vol.1.
- Besacier, L., Barnard, E., Karpov, A., Schultz, T., 2014. Automatic Speech Recognition for Under-resourced Languages: A Survey. *Speech Communication* 56, 85–100.
- Bisani, M., Ney, H., 2004. Bootstrap estimates for confidence intervals in ASR performance evaluation, in: Proc. of ICASSP, pp. I–409–12.
- Cetin, O., Kantor, A., King, S., Bartels, C., Magimai-Doss, M., Frankel, J., Livescu, K., 2007a. An Articulatory Feature-Based Tandem Approach and Factored Observation Modeling, in: Proc. of ICASSP, pp. 645–648.
- Cetin, O., Magimai-Doss, K., Livescu, K., Kantor, A., King, S., Bartels, C., Frankel, J., 2007b. Monolingual and crosslingual comparison of tandem features derived from articulatory and phone MLPs, in: Proc. of ASRU, pp. 36–41.
- Chang, S., 2002. A Syllable, Articulatory-Feature, and Stress-Accent Model of Speech Recognition. Ph.D. thesis. University of California, Berkeley.
- Chomsky, N., Halle, M., 1968. *The Sound Pattern of English*. MIT Press.
- Dahl, G., Yu, D., Deng, L., Acero, A., 2012. Context-Dependent Pre-Trained Deep Neural Networks for Large-Vocabulary Speech Recognition. *IEEE Trans. on Audio, Speech, and Language Processing* 20, 30–42.
- Deng, L., Ramsay, G., Sun, D., 1997. Production Models As a Structural Basis for Automatic Speech Recognition. *Speech Communication* 22, 93–111.
- Dines, J., Magimai-Doss, M., 2007. A Study of Phoneme and Grapheme based Context-Dependent ASR Systems, in: Proc. of Machine Learning for Multimodal Interaction (MLMI), pp. 215–226.
- Frankel, J., King, S., 2005. A Hybrid ANN/DBN Approach to Articulatory Feature Recognition, in: Proc. of Interspeech, pp. 3045–3048.
- Frankel, J., Wester, M., King, S., 2007. Articulatory feature recognition using dynamic Bayesian networks, in: *Computer Speech & Language*, pp. 620–640.
- Greenberg, S., Chang, S., Hollenback, J., 2000. An introduction to the diagnostic evaluation of Switchboard-corpus automatic speech recognition systems, in: *Proceedings of the NIST Speech Transcription Workshop*.
- Hain, T., 2005. Implicit Modelling of Pronunciation Variation in Automatic Speech Recognition. *Speech Communication* 46, 171–188.
- Hain, T., Woodland, P.C., 1999. Dynamic HMM Selection for Continuous Speech Recognition, in: Proc. of EUROSPEECH, pp. 1327–1330.
- Harris, J., 1994. *English Sound Structure*. Blackwell.

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9 Hermansky, H., Ellis, D., Sharma, S., 2000. Tandem connectionist feature extraction for
10 conventional HMM systems, in: Proc. of ICASSP, pp. 1635–1638.
- 11 Hinton, G., Deng, L., Yu, D., Dahl, G., Mohamed, A.R., Jaitly, N., Senior, A., Vanhoucke,
12 V., Nguyen, P., Sainath, T., Kingsbury, B., 2012. Deep neural networks for acoustic
13 modeling in speech recognition: The shared views of four research groups. IEEE
14 Signal Processing Magazine 29, 82–97.
- 15
- 16 Hiroya, S., Honda, M., 2004. Estimation of articulatory movements from speech acous-
17 tics using an HMM-based speech production model. IEEE Trans. on Speech and Audio
18 Processing 12, 175–185.
- 19 Hogden, J., Lofqvist, A., Gracco, V., Zlokarnik, I., Rubin, P., Saltzman, E., 1996. Accurate
20 recovery of articulator positions from acoustics: New conclusions based on human
21 data. The Journal of The Acoustical Society of America 100, 1819–1834.
- 22
- 23 Holter, T., Svendsen, T., 1997. Incorporating Linguistic Knowledge And Automatic Base-
24 form Generation In Acoustic Subword Unit Based Speech Recognition, in: Proc. of
25 EUROSPEECH.
- 26 Hosom, J., 2009. Speaker-independent phoneme alignment using transition-dependent
27 states. Speech Communication 51, 352–368.
- 28
- 29 Huang, X., Jack, M., 1989. Semi-continuous hidden Markov models for speech signal.
30 Computer Speech and Language 3, 239–251.
- 31
- 32 Huttenlocher, D., Zue, V., 1984. A model of lexical access from partial phonetic informa-
33 tion, in: Proc. of ICASSP, pp. 391–394.
- 34 Hwang, M.Y., Huang, X., 1992. Subphonetic modeling with Markov states-Senone, in:
35 Proc. of ICASSP, pp. 33–36 vol.1.
- 36
- 37 Imseng, D., Dines, J., Motlicek, P., Garner, P., , Bourlard, H., 2012. Comparing different
38 acoustic modeling techniques for multilingual boosting, in: Proc. of Interspeech.
- 39
- 40 Juneja, A., Espy-Wilson, C., 2004. Significance of invariant acoustic cues in a prob-
41 abilistic framework for landmark-based speech recognition, in: From Sound to
42 Sense: Fifty+ Years of Discoveries in Speech Communication, MIT, Cambridge, Mas-
43 sachusetts.
- 44 Jyothi, P., 2013. Discriminative And Articulatory Feature-Based Pronunciation Models
45 For Conversational Speech Recognition. Ph.D. thesis. The Ohio State University, Ohio.
- 46 Jyothi, P., Fosler-Lussier, E., Livescu, K., 2013. Discriminative Training of WFST Factors
47 with Application to Pronunciation Modeling, in: Proc. of Interspeech, pp. 1961–1965.
- 48
- 49 Jyothi, P., Livescu, K., Fosler-Lussier, E., 2011. Lexical access experiments with context-
50 dependent articulatory feature-based models, in: Proc. of ICASSP, pp. 4900–4903.
- 51
- 52 King, S., Frankel, J., Livescu, K., McDermott, E., Richmond, K., Wester, M., 2007. Speech
53 production knowledge in automatic speech recognition . The Journal of The Acoustical
54 Society of America 121, 723–742.
- 55
- 56 King, S., Taylor, P., 2000. Detection of Phonological Features in Continuous Speech using
57 Neural Networks. Computer Speech and Language 14, 333–353.
- 58
- 59
- 60
- 61
- 62
- 63
- 64
- 65

- Kirchhoff, K., 1996. Syllable-level desynchronisation of phonetic features for speech recognition, in: Proceedings of Int. Conf. Spoken Language Processing.
- Kirchhoff, K., Fink, G.A., Sagerer, G., 2002. Combining acoustic and articulatory feature information for robust speech recognition . *Speech Communication* 37, 303–319.
- Ladefoged, P., 1993. *A Course in Phonetics*. Harcourt Brace College Publishers.
- Lal, P., King, S., 2013. Cross-Lingual Automatic Speech Recognition Using Tandem Features. *IEEE Trans. on Audio, Speech, and Language Processing* 21, 2506–2515.
- Livescu, K., Cetin, O., Hasegawa-Johnson, M., King, S., Bartels, C., Borges, N., Kantor, A., Lal, P., Yung, L., Bezman, A., Dawson-Haggerty, S., Woods, B., 2008. Articulatory Feature-based Methods for Acoustic and Audio-Visual Speech Recognition: 2006 JHU Summer Workshop Final Report. Technical Report. John Hopkins University, Center for Language and Speech Processing.
- Livescu, K., Glass, J.R., 2004. Feature-based pronunciation modeling with trainable asynchrony probabilities, in: Proceedings of Int. Conf. Spoken Language Processing.
- Luo, X., Jelinek, F., 1999. Probabilistic Classification of HMM States for Large Vocabulary Continuous Speech Recognition, in: Proc. of ICASSP, pp. 353–356.
- Magimai-Doss, M., Rasipuram, R., Aradilla, G., Boulard, H., 2011. Grapheme-based Automatic Speech Recognition using KL-HMM, in: Proc. of Interspeech, pp. 2693–2696.
- Metze, F., Waibel, A., 2002. A flexible stream architecture for ASR using articulatory features, in: Proceedings of Int. Conf. Spoken Language Processing.
- Morgan, N., Boulard, H., 1995. Continuous speech recognition: An introduction to the hybrid HMM/connectionist approach. *IEEE Signal Processing Magazine* , 25–42.
- Næss, A.B., Livescu, K., Prabhavalkar, R., 2011. Articulatory Feature Classification Using Nearest Neighbors, in: Proc. of Interspeech, pp. 2301–2304.
- Paul, D., Baker, J., 1992. The Design for the Wall Street Journal-based CSR Corpus, in: DARPA Speech and Language Workshop, Morgan Kaufmann Publishers.
- Pinto, J., Sivaram, G.S.V.S., Magimai-Doss, M., Hermansky, H., Boulard, H., 2011. Analysis of MLP Based Hierarchical Phoneme Posterior Probability Estimator. *IEEE Trans. on Audio, Speech, and Language Processing* 19, 225–241.
- Povey, D., Burget, L., Agarwal, M., Akyazi, P., Kai, F., Ghoshal, A., Glembek, O., Goel, N., Karafiát, M., Rastrow, A., Rose, R., Schwarz, P., Thomas, S., 2011a. The sub-space Gaussian mixture model - A structured model for speech recognition. *Computer Speech and Language* .
- Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlicek, P., Qian, Y., Schwarz, P., Silovsky, J., Stemmer, G., Vesely, K., 2011b. The Kaldi Speech Recognition Toolkit, in: Proc. of ASRU.
- Prabhavalkar, R., Fosler-Lussier, E., Livescu, K., 2011. A factored conditional random field model for articulatory feature forced transcription, in: Proc. of ASRU, pp. 77–82.

- Price, P.J., Fisher, W., Bernstein, J., 1988. The DARPA 1000-word resource management database for continuous speech recognition, in: Proc. of ICASSP, pp. 651–654.
- Rabiner, L., 1989. A tutorial on hidden Markov models and selected applications in speech recognition. Proc. of IEEE 77, 257–286.
- Rasipuram, R., 2014. Grapheme-based Automatic Speech Recognition using Probabilistic Lexical Modeling. Ph.D. thesis. EPFL. Switzerland.
- Rasipuram, R., Magimai-Doss, M., 2011a. Improving articulatory feature and phoneme recognition using multitask learning, in: Artificial Neural Networks and Machine Learning - ICANN 2011.
- Rasipuram, R., Magimai-Doss, M., 2011b. Integrating articulatory features using Kullback-Leibler divergence based acoustic model for phoneme recognition, in: Proc. of ICASSP.
- Rasipuram, R., Magimai-Doss, M., 2011c. Multitask Learning to Improve Articulatory Feature Estimation and Phoneme Recognition. Idiap Research Report, Idiap-RR-21-2011.
- Rasipuram, R., Magimai-Doss, M., 2012. Acoustic Data-driven Grapheme-to-Phoneme Conversion using KL-HMM, in: Proc. of ICASSP.
- Rasipuram, R., Magimai-Doss, M., 2013a. Improving Grapheme-based ASR by Probabilistic Lexical Modeling Approach, in: Proceedings of Interspeech.
- Rasipuram, R., Magimai-Doss, M., 2013b. Probabilistic Lexical Modeling and Grapheme-based Automatic Speech Recognition. http://publications.idiap.ch/downloads/reports/2013/Rasipuram_Idiap-RR-15-2013.pdf . Idiap Research Report.
- Rasipuram, R., Magimai-Doss, M., 2015. Acoustic and lexical resource constrained ASR using language-independent acoustic model and language-dependent probabilistic lexical model. Speech Communication 68, 23–40. URL: <http://www.sciencedirect.com/science/article/pii/S0167639314000995>.
- Richardson, M., Bilmes, J., Diorio, C., 2003. Hidden-Articulatory Markov Models for Speech Recognition. Speech Communication 41, 713–716.
- Rottland, J., Rigoll, G., 2000. Tied Posteriors: An Approach for Effective Introduction of Context Dependency in Hybrid NN/HMM LVCSR, in: Proc. of ICASSP, pp. 1241–1244.
- Saraclar, M., Nock, H., Khudanpur, S., 2000. Pronunciation Modeling By Sharing Gaussian Densities Across Phonetic Models, in: Computer Speech and Language, pp. 137–160.
- Scharenborg, O., Wan, V., Moore, R.K., 2007. Towards capturing fine phonetic variation in speech using articulatory features. Speech Communication 49, 811–826.
- Siniscalchi, S., Lyu, D.C., Svendsen, T., Lee, C.H., 2012. Experiments on Cross-Language Attribute Detection and Phone Recognition With Minimal Target-Specific Training Data. IEEE Trans. on Audio, Speech, and Language Processing 20, 875–887.
- Strik, H., Cucchiaroni, C., 1999. Modeling pronunciation variation for ASR: A survey of the literature. Speech Communication 29, 225–246.

- Stüker, S., Metze, F., Schultz, T., Waibel, A., 2003. Integrating Multilingual Articulatory Features into Speech Recognition, in: Proc. of Eurospeech, pp. 1033–1036.
- Suzuki, S., Okadome, T., Honda, M., 1998. Determination of articulatory positions from speech acoustics by applying dynamic articulatory constraints, in: Proceedings of Int. Conf. Spoken Language Processing.
- Wester, M., Frankel, J., King, S., 2004. Asynchronous articulatory feature recognition using dynamic Bayesian networks, in: Proc. of IEICI Beyond HMM Workshop.
- Woodland, P.C., Odell, J.J., Valtchev, V., Young, S.J., 1994. Large Vocabulary Continuous Speech Recognition using HTK, in: Proc. of ICASSP, pp. 125–128.
- Young, S.J., Odell, J.J., Woodland, P.C., 1994. Tree-based State Tying for High Accuracy Acoustic Modelling, in: Proceedings of the Workshop on Human Language Technology (HLT), pp. 307–312.

Appendix A. Parameter Estimation in the KL-HMM Approach

Given a trained ANN and a training set of N utterances $\{X(n), W(n)\}_{n=1}^N$, the set of acoustic unit probability vectors $\{Z(n), W(n)\}_{n=1}^N$ are estimated. For each training utterance n , $X(n)$ represents the sequence of cepstral features, $W(n)$ represents the sequence of underlying words, $Z(n)$ represents a sequence of acoustic unit probability vectors and $T(n)$ represents the number of cepstral features or the number of acoustic unit probability vectors.

The KL-HMM system is parameterized by $\Theta_{\text{kuil}} = \{\{y_i\}_{i=1}^I, \{a_{ij}\}_{i,j=1}^I\}$. The lexical model parameters $\{y_i\}_{i=1}^I$ are initialized uniformly, i.e., initially $y_i^d = \frac{1}{D} \forall i, d$. The training data $\{Z(n), W(n)\}_{n=1}^N$ and the current parameter set Θ_{kuil} , are used to estimate the new set of parameters $\hat{\Theta}_{\text{kuil}}$ by the Viterbi algorithm. In the case of the local score S_{RKL} the cost function minimized is,

$$\hat{\Theta}_{\text{kuil}} = \arg \min_{\Theta_{\text{kuil}}} \left[\sum_{n=1}^N \min_{Q \in \mathcal{Q}} \sum_{t=1}^{T(n)} [S_{\text{RKL}}(\mathbf{y}_{q_t}, \mathbf{z}_t(n)) - \log a_{q_{t-1}q_t}] \right] \quad (\text{A.1})$$

where $Q = \{q_1, \dots, q_t, \dots, q_{T(n)}\}$, $q_t \in \{1, \dots, I\}$ and $\mathcal{Q}$ denotes set of all possible HMM state sequences.

The training process involves iteration over the segmentation and the optimization steps until convergence. Given the current set of parameters, the segmentation step yields an optimal state sequence for each training utterance using the Viterbi algorithm. Given optimal state sequences and acoustic unit posterior vectors belonging to the states, the optimization step then estimates new set of model parameters by minimizing Eqn. (A.1) subject to the constraint that $\sum_{d=1}^D y_i^d = 1$.

The optimal state distribution for the local score S_{RKL} (Equation (16)), is the arithmetic mean of the training acoustic unit probability vectors assigned to the state, i.e.,

$$y_i^d = \frac{1}{M(i)} \sum_{\mathbf{z}_t(n) \in Z(i)} z_t^d(n) \quad \forall d \quad (\text{A.2})$$

where $Z(i)$ denotes the set of acoustic unit probability vectors assigned to state i and $M(i)$ is the cardinality of $Z(i)$. More details about the parameter estimation for the KL-HMM approach can be found in the thesis by Aradilla (2008).

Appendix B. Details of the Experimental Study

Appendix B.1. Phoneme-to-Articulatory Feature Map

Table B.1: Knowledge-based phoneme-to-articulatory feature map used in this paper

Phoneme	Manner	Place	Height	Vowel
sil	sil	sil	sil	sil
aa	vowel	back	low	aa
ae	vowel	mid-front	low	ae
ah	vowel	mid	mid	ah
ao	vowel	back	mid-low	ao
aw1	vowel	mid-front	low	aw1
aw2	vowel	mid-back	high	aw2
ax	vowel	mid	mid	ax
axr	approximant	retroflex	mid	consonant
ay1	vowel	back	low	ay1
ay2	vowel	mid-front	high	ay2
b	voiced-stop	labial	max	consonant
ch	stop	front	max	consonant
d	voiced-stop	alveolar	max	consonant
dh	voiced-fricative	dental	max	consonant
eh	vowel	mid-front	mid	eh
el	approximant	lateral	very-high	consonant
em	nasal	labial	max	consonant
en	nasal	alveolar	max	consonant
er	vowel	mid	mid	er
ey1	vowel	front	mid-high	ey1
ey2	vowel	mid-front	high	ey2
f	fricative	labial	max	consonant
g	voiced-stop	dorsal	max	consonant
hh	aspirated	unknown	max	consonant
ih	vowel	mid-front	high	ih
iy	vowel	front	very-high	iy
jh	voiced-stop	front	max	consonant
k	stop	dorsal	max	consonant
l	approximant	lateral	very-high	consonant
m	nasal	labial	max	consonant
n	nasal	alveolar	max	consonant
ng	nasal	dorsal	max	consonant
ow1	vowel	back	mid	ow1
ow2	vowel	mid-back	high	ow2
oy1	vowel	back	mid-low	oy1
oy2	vowel	mid-front	high	oy2
p	stop	labial	max	consonant
r	approximant	retroflex	mid-low	consonant
s	fricative	alveolar	max	consonant
sh	fricative	front	max	consonant
t	stop	alveolar	max	consonant
th	fricative	dental	max	consonant
uh	vowel	mid-back	high	uh
uw	vowel	back	very-high	uw
v	voiced-fricative	labial	max	consonant
w	approximant	back	very-high	consonant
y	approximant	front	very-high	consonant
z	voiced-fricative	alveolar	max	consonant
zh	voiced-fricative	front	max	consonant

Table B.2: Overview of the various MLPs used in the paper. The number of input, output and hidden units of MLPs are denoted as i , h and o . The total number of parameters of an MLP is computed as $h * (i + o + 1)$. The frame accuracies of the MLPs *MLP-RM-PH*, *MLP-RM-AF*, *MULTI-WSJRM-PH* and *MULTI-WSJRM-AF* are computed on the development set of the RM corpus. The frame accuracies of the MLPs *MLP-WSJ-PH*, *MLP-WSJ-AF* are computed on the development set of the WSJ corpus.

MLP	acoustic units	input size (i)	hidden size (h)	output size (o)	MLP size	frame acc.
<i>MLP-RM-PH</i>	phonemes	$39*9=351$	1260	45	0.5M	73.77
<i>MLP-WSJ-PH</i>	phonemes	$39*9=351$	3652	45	1.5M	69.34
<i>MULTI-WSJRM-PH</i>	phonemes	$17*45=765$	730	45	2.1M	80.16
<i>MLP-RM-AF</i>	manner	$39*9=351$	682	10	0.2M	82.48
	place	$39*9=351$	676	13	0.2M	76.20
	height	$39*9=351$	685	8	0.2M	79.26
	vowel	$39*9=351$	658	23	0.2M	79.38
<i>MLP-WSJ-AF</i>	manner	$39*9=351$	4005	10	1.4M	89.18
	place	$39*9=351$	3972	13	1.4M	86.83
	height	$39*9=351$	4027	8	1.4M	88.10
	vowel	$39*9=351$	3866	23	1.4M	88.29
<i>MULTI-WSJRM-AF</i>	manner	$17*54=918$	1063	10	6.6M	86.47
	place	$17*54=918$	1059	13	6.6M	82.42
	height	$17*54=918$	1065	8	6.6M	84.25
	vowel	$17*54=918$	1048	23	6.6M	84.36

Appendix B.2. Details of the MLPs

The number of hidden units of the MLPs are selected based on the optimal frame accuracy on the development set of the RM or WSJ corpora. Hence, the total number of parameters are different for various MLPs. For the AF-based MLPs trained on the RM corpus, optimal frame accuracy was obtained when the total number of MLP parameters were about 15% of the number of training frames. For the phoneme-based and AF-based MLPs trained on the WSJ corpus, it was observed that the optimal frame accuracy on the development set was obtained when the total number of parameters were about 5% of the number of training frames (28M). Therefore, the total number of parameters of the MLPs trained on the WSJ corpus is about $28M * 0.05 \approx 1.5M$. The total number of parameters in the multistage MLPs consist of the parameters of the first stage MLP(s) and the second stage MLP.

Appendix B.3. Details of Various Systems

Table B.3: The details of the various KL-HMM systems, namely the total number of MLP parameters (N_{θ_a}), lexical units (I), acoustic units (D), lexical model parameters (N_{θ_l}), KL-HMM system parameters ($N_{\theta_a} + N_{\theta_l}$). The details of the various tandem systems, namely the number tied states (ts), Gaussian components in each tied state (ts), dimensionality of the tandem features that contribute to 99% of the variance (v), tandem system parameters excluding the MLP parameters ($N_{\theta_{tan}}$) and tandem system parameters with the MLP parameters ($N_{\theta_a} + N_{\theta_{tan}}$). The total number of parameters in the cross-word context-dependent phoneme-based and grapheme-based HMM/GMM systems are about 1.0M.

MLP	N _{θ_a}	KL-HMM system				Tandem system			
		#lexical units (l)	#acoustic units (D)	N _{θ_l} I * D	Total N _{θ_a} + N _{θ_l}	#tied states (ts), #Gaussians (g)	dim (v)	Tandem (N _{θ_{a+n}}) ts * (g * (2v + 1))	Total N _{θ_a} + N _{θ_{a+n}}
MLP-RM-PH	0.5M	10988	45	0.5M	1.0M	1726,8	39	1.1M	1.6M
MLP-RM-AF	0.8M	8127	54	0.4M	1.2M	1737,8	42	1.2M	2.0M
MLP-RM-PH+AF	1.3M	11678	99	1.1M	2.4M	1723,8	75	2.1M	3.4M
MLP-WSJ-PH	1.5M	11723	45	0.5M	2.0M	2367,8	37	1.4M	2.9M
MLP-WSJ-AF	5.6M	8059	54	0.4M	6.0M	1751,8	42	1.2M	6.8M
MLP-WSJ-PH+AF	7.1M	9753	99	1.0M	8.1M	1718,8	75	2.1M	9.2M
MULTI-WSJRM-PH	2.1M	10857	45	0.5M	2.6M	2546,8	36	1.5M	3.6M
MULTI-WSJRM-AF	9.6M	7753	54	0.4M	10.0M	1807,8	40	1.2M	10.8M
MULTI-WSJRM-PH+AF	11.7M	9975	99	1.0M	12.7M	1784,8	73	2.1M	13.8M