LANGUAGE RECOGNITION WITH DISCRIMINATIVE KEYWORD SELECTION*

F. S. Richardson, W. M. Campbell

MIT Lincoln Laboratory

ABSTRACT

One commonly used approach for language recognition is to convert the input speech into a sequence of tokens such as words or phones and then to use these token sequences to determine the target language. The language classification is typically performed by extracting N-gram statistics from the token sequences and then using an N-gram language model or support vector machine (SVM) to perform the classification. One problem with these approaches is that the number of N-grams grows exponentially as the order N is increased. This is especially problematic for an SVM classifier as each utterance is represented as a distinct N-gram vector. In this paper we propose a novel approach for modeling higher order N-grams using an SVM via an alternating filter-wrapper feature selection method. We demonstrate the effectiveness of this technique on the NIST 2007 language recognition task.

Index Terms— Language Recognition, Support Vector Machines

1. INTRODUCTION

Language recognition systems generally fall into two categories: acoustic systems which model short term spectral characteristics of speech and token systems which model *high level* features such as words or phones. State-of-the-art systems typically include both techniques to achieve optimal performance. In this paper, we focus on a token based language recognition system which uses a phone recognizer for generating the token sequences and a support vector machine for classifying the target language.

Token-based SVM systems have been used in many speaker and language recognition systems [1, 2, 3]. Sequence kernels are constructed by viewing a speech segment as a *document* of tokens. The SVM feature space in this case is a scaling of co-occurrence probabilities (or *N*-grams) of tokens in an utterance. This technique is analogous to methods for applying SVMs to text classification [4].

In our prior work [5], we used the output of a speech recognition system to generate sequences of words. We observed that certain words (or keywords) were strongly indicative of a specific language. For example, the phrase "you know" was a strong indicator that the language was English. One limitation with this approach was that the possible set of words for distinguishing between languages was limited to the word lexicon of the underlying speech recognition system.

In this paper, our token sequences are generated using a phone recognizer. We then attempt to find keywords by selecting the N-grams which discriminate the most between the different languages for progressively higher orders of N. Unfortunately, this problem is

difficult to solve directly, since the number of unique N-grams grows exponentially as we increase the order. To alleviate this difficulty, we propose a method that starts by selecting a relatively small set of lower order N-grams and then using these sets to successively build candidate lists of higher order N-grams.

The outline of the paper is as follows. In Section 2.1, we review the basic architecture that we use for phone recognition and how it is applied to the problem. In Section 2.2, we review the application of SVMs to determining a target language. Section 3.1 describes a feature selection method for SVMs. Section 3.2 presents our method for constructing long context units of phones to automatically create keywords. We use a novel feature selection approach that attempts to find longer strings that discriminate well between classes. Finally, in Section 4, we show the application of our method to a language recognition problem. We show qualitatively that the method produces interesting keywords. Quantitatively, we show that the method produces keywords which are good discriminators between classes.

2. PHONOTACTIC CLASSIFICATION

2.1. Phone Recognition

The high-level component of our system is a phone recognizer based upon the Brno University (BUT) design [6]. The basic architecture of this system is a three state monophone HMM system with a null grammar. This system uses two powerful components to achieve high accuracy. First, split temporal context (STC) features provide contextual cues for modeling monophones. Second, the BUT recognizer makes extensive use of discriminatively trained feedforward artificial neural nets to model HMM state posterior probabilities.

We developed an English phone recognizer using the BUT architecture and automatically generated STT transcripts on the Switchboard 2 Cell corpora [7]. Training data consisted of approximately 10 hours of speech. ANN training was accomplished using the ICSI Quicknet package [8]. The resulting system has 49 monophones including silence.

The BUT recognizer is used along with the HTK HMM toolkit [9] to produce lattices. Lattices encode multiple hypotheses with acoustic likelihoods. We use the lattice to produce expected counts of tokens and token *N*-grams.

Expected counts of N-grams can be easily understood as an extension of standard counts. Suppose we have a hypothesized string of tokens, $W = w_1, \dots, w_n$. Then bigrams are created by grouping two tokens at a time to form, $W_2 = w_1 _ w_2, w_2 _ w_3, \dots, w_{n-1} _ w_n$. Higher order N-grams are formed from longer juxtapositions of tokens. The count function for a given bigram, d_i , is count $(d_i|W_2)$ is the number of occurrences of d_i in the sequence W_2 . To extend counts to a lattice, \mathcal{L} , we find the expected count over all all possible

^{*}This work was sponsored by the Department of Defense under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

hypotheses in the lattice,

$$\operatorname{count}(d_i|\mathcal{L}) = E_W[\operatorname{count}(d_i|W)] = \sum_{W \in \mathcal{L}} p(W|\mathcal{L})\operatorname{count}(d_i|W).$$
(1)

The expected counts can be computed efficiently by a forwardbackward algorithm; more details can be found in Section 3.3 and [10].

A useful application of expected counts is to find the probability of an *N*-gram in a lattice. For a lattice, \mathcal{L} , the joint probability of an *N*-gram, d_i , is

$$p(d_i|\mathcal{L}) = \frac{\operatorname{count}(d_i|\mathcal{L})}{\sum_j \operatorname{count}(d_j|\mathcal{L})}$$
(2)

where the sum in (2) is performed over all *unique* N-grams in the utterance. Note that the denominator in (2) is approximately constant as N varies.

2.2. Discriminative Language Modeling: SVMs

Token-based class recognition using SVMs can be performed in several ways [11, 3]. We focus on an approach which is similar to [11, 2]. In [2], a token sequence, $W = w_1, \dots, w_N$, is modeled using a bag-of-*N*-grams. For a sequence of tokens, (joint) probabilities of the unique *N*-grams, d_j , on a per conversation basis are calculated, $p(d_j|W)$. Then, the probabilities are mapped to a sparse vector with entries

$$D_j p(d_j | W). \tag{3}$$

The selection of the weighting, D_j , in (3) is critical for good performance. A typical choice has form

$$D_j = \min\left(C_j, g_j\left(\frac{1}{p(d_j|\text{all})}\right)\right) \tag{4}$$

where $g_j(\cdot)$ is a function which squashes the dynamic range, and C_j is a constant. The probability $p(d_j|\text{all})$ in (4) is calculated from the observed probability across all classes. We choose $g_j(x) = \sqrt{x}$ and $C_j = \infty$.

The general weighting of probabilities is then combined to form a kernel between two lattices, see [2] for more details. For two lattices, \mathcal{L}_1 and \mathcal{L}_2 , the kernel is

$$K(\mathcal{L}_1, \mathcal{L}_2) = \sum_j D_j^2 p(d_j | \mathcal{L}_1) p(d_j | \mathcal{L}_2).$$
(5)

Intuitively, the kernel in (5) says that if the same *N*-grams are present in two sequences and the normalized frequencies are similar there will be a high degree of similarity (a large inner product). If *N*grams are not present, then this will reduce similarity since one of the probabilities in (5) will be zero. The normalization D_j insures that *N*-grams with large probabilities do not dominate the kernel function. The kernel can alternatively be viewed as a linearization of the log-likelihood ratio [2].

Incorporating the kernel (5) into an SVM system is straightforward. SVM training and scoring require only a method of kernel evaluation between two objects that produces positive definite kernel matrices (the Mercer condition). We use the package SVM-Torch [12]. Training is performed with a one-versus-all strategy. For each target class, we group all remaining class data and then train with these two classes.

3. DISCRIMINATIVE KEYWORD SELECTION

3.1. SVM Feature Selection

A first step towards an algorithm for automatic keyword generation using phones is to examine feature selection methods. Ideally, we would like to select over all possible N-grams, where N is varying, the most discriminative sequences for determining the language of a speech segment. The number of features in this case is prohibitive, since it grows exponentially with N. Therefore, we have to consider alternate methods.

As a first step, we examine feature selection for fixed N and look for keywords with N or less phones. Suppose that we have a set of candidate keywords. Since we are already using an SVM, a natural algorithm for discriminative feature selection in this case is to use a wrapper method [13].

Suppose that the optimized SVM solution is

$$f(X) = \sum_{i} \alpha_i K(X, X_i) + c \tag{6}$$

and

$$\mathbf{w} = \sum_{i} \alpha_i \mathbf{b}(X_i) \tag{7}$$

where $\mathbf{b}(X_i)$ is the vector of weighted *N*-gram probabilities in (3). We note that the kernel given in (5) is linear. Also, the *N*-gram probabilities have been normalized in (3) by their probability across the entire data set. Intuitively, because of this normalization and since $f(X) = \mathbf{w}^t \mathbf{b}(X) + c$, large magnitude entries in \mathbf{w} correspond to significant features.

A confirmation of this intuitive idea is the algorithm of Guyon, et. al. [14]. Guyon proposes an iterative wrapper method for feature selection for SVMs which has these basic steps:

- For a set of features, S, find the SVM solution with model w.
- Rank the features by their corresponding model entries w²_i. Here, w_i is the *i*th entry of w in (7).
- Eliminate low ranking features.

The algorithm may be iterated multiple times.

Guyon's algorithm for feature selection can be used for picking significant N-grams as keywords. We can create a kernel which is the sum of kernels as in (5) up to the desired N. We then train an SVM and rank N-grams according to the magnitude of the entries in the SVM model vector, \mathbf{w} .

As an example, we have looked at this feature selection method for a language recognition task with trigrams (to be described in Section 4). As a motivation for the applicability of Guyon's feature selection method, see Figure 1. The figure shows two functions. First, the CDF of the SVM model values, $|w_i|$, is shown. The CDF has an S-curve shape indicating that there are a large set of small model weights, and a small set of large model weights.

The second curve shows the equal error rate (EER) of the task as a function of applying one iteration of the Guyon algorithm and retraining the SVM. EER is defined as the value where the miss and false alarm rates are equal. All features with $|w_i|$ below the value on the x-axis are discarded in the first iteration. From the figure, we see that only a small fraction (< 5%) of the features are needed to obtain good error rates. This interesting result provides motivation that a small subset of keywords are significant to the task.



Fig. 1. Feature selection for a trigram language recognition task using Guyon's method

3.2. Keywords via an alternating wrapper/filter method

The algorithm in Section 3.1 gives a method for N-gram selection for fixed N. Now, suppose we want to find keywords for arbitrary N. One possible hypothesis for keyword selection is that since higher order N-grams are discriminative, lower order N-grams in the keywords will also be discriminative. Therefore, it makes sense to find distinguishing lower order N-grams and then construct longer units from these. On the basis of this idea, we propose the following algorithm for keyword construction:

Keyword Building Algorithm

- Start with an initial value of N = N_s. Initialize the set, S'_N, to all possible N-grams of phones for this N. For example, S'₃ would be the set of all possible trigrams (this is a relatively small set for a phone decoder).
- Wrapper Step Apply the feature selection algorithm in Section 3.1 to produce a subset of distinguishing N-grams, S_N ⊂ S'_N.
- Filter Step Construct a new set of (N + 1)-grams by juxtaposing elements from S_N with phones. Nominally, we take this step to be juxtaposition on the left and right, $S'_{n+1} = \{pd, dq | d \in S_n, p \in S_1, q \in S_1\}.$
- Set $N \leftarrow N+1,$ iterate the wrapper and filter steps to some desired order N
- Output: S_N

A few items should be noted about the proposed keyword building algorithm. First, we call the second feature selection process a filter step, since induction has not been applied to produce the (N + 1)-gram features. Second, note that the purpose of the filter step is to provide a candidate set of possible (N + 1)-grams which can then be more systematically reduced. Third, several potential algorithms exist for the filter step. In our experiments and in the algorithm description, we nominally append one phone to the beginning and end of an N-gram. Another possibility is to try to combine overlapping N-grams. For instance, suppose the keyword is some_people which has the phone transcript s_ah_m_p_iy_p_l. Then, if we are looking at 4-grams, we might see as top features s_ah_m_p and p_iy_p_l and combine these to produce a new keyword.

3.3. Keyword Implementation

The expected N-gram counts were computed from lattices using the forward-backward algorithm. Equation (8) gives the posterior probability of a connected sequence of arcs in the lattice where $src_nd(a)$ and $dst_nd(a)$ are the source and destination node of arc a, $\ell(a)$ is the likelihood associated with arc a, $\alpha(n)$ and $\beta(n)$ are the forward and backward probabilities of reaching node N from the beginning or end of the lattice \mathcal{L} respectively, and $\ell(\mathcal{L})$ is the total likelihood of the lattice (the $\alpha(\cdot)$ of the final node or $\beta(\cdot)$ of the initial node of the lattice).

$$p(a_j, \dots, a_{j+n}) = \frac{\alpha(\operatorname{src_nd}(a_j))\ell(a_j) \dots \ell(a_{j+n})\beta(\operatorname{dst_nd}(a_{j+n}))}{\ell(\mathcal{L})}$$
(8)

Now if we define the posterior probability of a node p(n) as $p(n) = (\alpha(n)\beta(n))/\ell(\mathcal{L})$. Then equation (8) becomes:

$$p(a_j,\ldots,a_{j+n}) = \frac{p(a_j)\ldots p(a_{j+n})}{p(src_nd(a_{j+1}))\ldots p(src_nd(a_{j+n}))}.$$
 (9)

Equation (9) is attractive because it provides a way of computing the path posteriors locally using only the individual arc and node posteriors along the path. We use this computation along with a *trie* structure [15] to compute the posteriors of our keywords.

4. EXPERIMENTS

4.1. Experimental setup

The phone recognizer described in Section 2.1 was used to generate lattices across a train and an evaluation data set. The training data set consists of more than 360 hours of telephone speech spanning 14 different languages and coming from a variety of different sources including Callhome, Callfriend, Fisher, Mixer and NIST 2007 evaluation development data. The evaluation data set is the NIST 2007 Language Recognition Evaluation data consisting of 7530 utterances spanning the 3, 10 and 30 second conditions. We evaluated our system for both the 30 and 10 second tasks under the NIST 2007 closed condition which limits the evaluation data to 14 languages (Arabic, Bengali, Chinese, English, Farsi, German, Hindustani, Japanese, Korean, Russian, Spanish, Tamil, Thai and vietnamese). These are the same conditions as the primary evaluation task for the NIST 2007 evaluation.

The training and evaluation data was segmented using an automatic speech activity detector and segments smaller than 0.5 seconds were thrown out. We also sub-segmented long audio files in the training data to keep the duration of each utterance to around 5 minutes (a shorter duration would have created too many training instances). For the wrapper step described in Section 3.2, the top and bottom 600 ranking keywords for each target language in the training data were selected according to the weights in each model. After the filter step, this resulted in a total of 785,097 keywords for the 4-gram, and a total of 970,890 keywords for the 5-gram. The support vector machine was trained using a dual formulation which requires pre-computing all of the kernel inner products between the data points and using an alternate kernel which simply indexes into the resulting kernel matrix. This approach becomes difficult when the number of data points is too large, but is good for large feature sets and parallelization.

4.2. Results (Qualitative and Quantitative)

To get a sense of how well our feature selection approach was working, we looked at the top ranking keywords from the English model

rank	phones	keyword		
1	SIL Y UW N OW	<sil> you know</sil>		
3	L AY K K SIL	like <sil></sil>		
12	!NULL SIL Y EH AX	<s> <sil> yeah</sil></s>		
14	IY HH AE AE V	I have		
16	SIL AE AE AE N	???		
19	SIL AE AE N D	and		
24	AA R R SIL !NULL	are <sil> </sil>		
26	SIL AY D OW N	<sil> I don't</sil>		
29	R AY AY T SIL	right <sil></sil>		
37	P IY P AX L	people		
48	DH AE T T SIL	that		

Table 1. Top ranking keywords for 5-gram SVM

N	1	2	3	4	5
10sec	21.2	12.0	07.8	06.6	09.2
30sec	13.4	04.3	02.2	01.8	03.1

 Table 2.
 English phone recognizer %EER for 10 and 30 second

 NIST 2007 tasks

only (since our phone recognizer is trained using the English phone set). Table 1 summarizes a few of the more compelling phone Ngrams (for 3, 4 and 5 grams), and a possible keyword that it corresponds to each one. Not surprisingly, we also noticed that in the list of top-ranking key words there were many variations and partial N-gram matches to what appeared to be the same keyword.

The equal error rate for our system on the NIST 2007 language recognition evaluation is summarized in Table 2. The 4-gram system gave a relative improvement of 15% on the 10 second task and 18% on the 30 second task, but despite the compelling keywords produced by the 5-gram system, the performance actually degraded significantly compared to the trigram and 4-gram systems. A comparison of using BUTs Hungarian tokenizer to our English tokenizer is presented in Table 3 along with results using a standard trigram language model (a standard 4-gram language model was not feasible for us to run). We've also demonstrated a significant reduction in EER by linearly fusing the English and Hungarian SVM scores for both the trigram and the 4-gram. A comparable fusion system using the standard trigram language model fused with a backend trained on held-out dev data does not perform nearly as well. Finally, we've presented the results for fusing the English and Hungarian 4-gram SVM systems with the GMM/LFA acoustic system used in the 2007 NIST evaluation. These two systems fused together perform within 10% of the MIT/LL evaluation system results.

5. CONCLUSIONS AND FUTURE WORK

We presented a method for automatic construction of keywords given a discriminative speech classification task. Our method was based upon an successively building longer span keywords from shorter span keywords using phones as a fundamental unit. The problem was cast as a feature selection method and an alternating filter and wrapper method was proposed. Results showed that reasonable keywords and improved performance could be achieved using this methodology.

6. REFERENCES

[1] Bin Ma and Haizhou Li, "A phonotactic-semantic paradigm for automatic spoken document classification," in *The 28th Annual Interna-*

System	N	30sec	10sec
EN SVM	3	2.2	7.8
EN SVM	4	1.8	6.6
HU SVM	3	2.6	9.0
HU SVM	4	2.4	8.1
Fused EN and HU SVM	3	1.6	5.8
Fused EN and HU SVM	4	1.5	5.2
Fused EN and GMM/LFA	4	1.1	4.0
Fused HU and GMM/LFA	4	1.4	4.3
MIT/LL Eval07 System	N/A	1.0	3.6
EN LM	3	3.5	8.5
HU LM	3	3.4	9.3
Fused EN and HU LM	3	2.3	7.1

Table 3. Hungarian and English SVM linear fusion compared to standard LM with backend fusion %EER for 10 and 30 second NIST tasks

tional ACM SIGIR Conference, Brazil, 2005.

- [2] W. M. Campbell, J. P. Campbell, D. A. Reynolds, D. A. Jones, and T. R. Leek, "Phonetic speaker recognition with support vector machines," in *Advances in Neural Information Processing Systems 16*, Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, Eds. MIT Press, Cambridge, MA, 2004.
- [3] W. M. Campbell, T. Gleason, J. Navratil, D. Reynolds, W. Shen, E. Singer, and P. Torres-Carrasquillo, "Advanced language recognition using cepstra and phonotactics: MITLL system performance on the NIST 2005 language recognition evaluation," in *Proc. IEEE Odyssey*, 2006.
- [4] T. Joachims, *Learning to Classify Text Using Support Vector Machines*, Kluwer Academic Publishers, 2002.
- [5] W. M. Campbell, F. Richardson, and D. A. Reynolds, "Language recognition with word lattices and support vector machines," in *Proceedings* of *ICASSP*, 2007, pp. IV–989 – IV–992.
- [6] Petr Schwarz, Matejka Pavel, and Jan Cernocky, "Hierarchical structures of neural networks for phoneme recognition," in *Proceedings of ICASSP*, 2006, pp. 325–328.
- [7] Linguistic Data Consortium, "Switchboard-2 corpora," http://www.ldc.upenn.edu.
- [8] "ICSI QuickNet," http://www.icsi.berkeley.edu/Speech/qn.html.
- [9] S. Young, Gunnar Evermann, Thomas Hain, D. Kershaw, Gareth Moore, J. Odell, D. Ollason, V. Valtchev, and P. Woodland, *The HTK book*, Entropic, Ltd., Cambridge, UK, 2002.
- [10] L. Rabiner and B.-H. Juang, *Fundamentals of Speech Recognition*, Prentice-Hall, 1993.
- [11] Lu-Feng Zhai, Man hung Siu, Xi Yang, and Herbert Gish, "Discriminatively trained language models using support vector machines for language identification," in *Proc. IEEE Odyssey: The Speaker and Language Recognition Workshop*, 2006.
- [12] Ronan Collobert and Samy Bengio, "SVMTorch: Support vector machines for large-scale regression problems," *Journal of Machine Learning Research*, vol. 1, pp. 143–160, 2001.
- [13] Avrim L. Blum and Pat Langley, "Selection of relevant features and examples in machine learning," *Artificial Intelligence*, vol. 97, no. 1-2, pp. 245–271, Dec. 1997.
- [14] I. Guyon, J. Weston, S. Barnhill, and V. Vapnik, "Gene selection for cancer classification using support vector machines," *Machine Learning*, vol. 46, no. 1-3, pp. 389–422, 2002.
- [15] Konrad Rieck and Pavel Laskov, "Language models for detection of unknown attacks in network traffic," *Journal of Computer Virology*, vol. 2, no. 4, pp. 243–256, 2007.