

Spotting Subsequences matching a HMM using the Average Observation Probability Criteria with application to Keyword Spotting

Marius Călin Silaghi

Florida Institute of Technology
msilaghi@fit.edu

Abstract

This paper addresses the problem of detecting keywords in unconstrained speech. The proposed algorithms search for the speech segment maximizing the average observation probability¹ along the most likely path in the hypothesized keyword model. As known, this approach (sometimes referred to as sliding model method) requires a relaxation of the begin/endpoints of the Viterbi matching, as well as a time normalization of the resulting score. This makes solutions complex (i.e., $\frac{LN^2}{2}$ basic operations for keyword HMM models with L states and utterances with N frames).

We present here two alternative (quite simple and efficient) solutions to this problem. a) First we provide a method that finds the optimal segmentation according to the criteria of maximizing the average observation probability. It uses Dynamic Programming as a step, but does not require scoring for all possible begin/endpoints. While the worst case remains $O(LN^2)$, this technique converged in at most $3(L+2)N$ basic operations in each experiment for two very different applications. b) The second proposed algorithm does not provide a segmentation but can be used for the decision problem of whether the utterance should be classified as containing the keyword or not (provided a predefined threshold on the acceptable average observation probability). This allows the algorithm to be even faster, with fix cost of $(L+2)N$.

Introduction

This paper addresses the problem of *keyword spotting (KWS)* in unconstrained speech. Typical applications are search in audio databases and recognition of spoken passwords, but due to their general framework the proposed algorithms also apply to computer vision, natural language processing, and bioinformatics (Silaghi 1999). Detecting the beginning and the end of the keyword segment in the utterance is called *segmentation*. This is typically done by explicit modeling of non-keyword segments using a filler Hidden Markov Model (HMM) or an ergodic HMM composed of context dependent or independent phone models without lexical constraints (Wilpon *et al.* 1990). The utterance is then modeled with HMMs where the keyword is preceded and followed by

non-keyword segments. The use of such non-keyword models, trained from speech databases, offers no guarantee that the hypothesized keyword segment maximizes any rigorous measure of the matching with the keyword model. Some slower approaches to KWS, without explicit modeling of non-keyword segments, have a stronger theoretical foundation. Namely, they maximize a chosen measure of the match between the keyword segment and a keyword model.

Although several algorithms tackling KWS without models of non-keyword segments have already been proposed in the past, (e.g., by using Dynamic Time Warping (Bridle 1973) or Viterbi matching (Wilpon *et al.* 1989; Boite *et al.* 1993; Junkawitsch *et al.* 1996; Benayed *et al.* 2003) allowing relaxation of the begin and endpoint constraints), these are known to require the use of an “appropriate” normalization of the matching scores since segments of different lengths have then to be compared. At any possible ending time e , the match score of the best warp and start time b of the reference has to be computed (Bridle 1973) (for all possible start times b associated with unpruned paths). Moreover, in (Wilpon *et al.* 1989), and in the same spirit than what is presented here, for all possible ending times e , the average observation likelihood along the most likely state sequence is used as scoring criterion. Sometimes these techniques are referred to as “sliding model methods” and need $\frac{LN^2}{2}$ updates to elements of a $L \times N$ Dynamic Programming (DP) phone trellis with L HMM states per keyword and utterances with N time frames. Previous attempts to construe faster alternatives based on straightforward DP (Junkawitsch *et al.* 1996) use cost functions that do not respect DP’s optimality principle (Bellman 1952) and are therefore sub-optimal from the point of view of the pursued measures.

Here we will use a similar rigorous scoring technique for keyword spotting like the one used by algorithms without filler models. We use an explicit filler model, re-estimated at each utterance in a way proven to yield optimal segmentation (for the chosen measure). Our matching score is based on the average observation probability (AOP), working with either likelihoods or posteriors. The first algorithm proposed here, besides computing the AOP score for the occurrence of the hypothesized keyword into an utterance, also produces the corresponding segmentation – the decision of where the keyword starts and ends. Compared to previously devised “sliding model” methods (such as (Bridle 1973;

Copyright © 2005, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

¹The accumulated log posterior or likelihood divided by the length of the segment.

Wilpon *et al.* 1989)), the algorithms proposed here re-estimate a filler model, using Viterbi decoding which does not require scoring for all begin/endpoints, and which can be proved to (quickly) converge to the “optimal solution” (from the point of view of the chosen scoring function, i.e., AOP). Each re-estimation step uses straightforward Viterbi alignments (similar to regular filler-based KWS).

Measures and Notations




Figure 1: A simple progressive left-to-right HMM.

A continuous Hidden Markov Model (HMM) (Jelinek 1999; Russell & Norvig 2002), see Figure 1, is defined by a set of states $\mathcal{Q} = \{q_1, \dots, q_L\}$ with a unique starting state, a space of possible outputs (for speech recognition, typically the space of Mel-frequency cepstrum coefficients feature vectors), a set of transition probabilities $P(q_i|q_j)$, $\forall i, j \in [1..L]$, and a set of probability density functions defining the likelihood of each possible output for each state.

Let $X = \{x_1, \dots, x_N\}$ denote the sequence of acoustic vectors in which we want to detect a keyword, and let M be the HMM model of a keyword and consisting of L states $\mathcal{Q} = \{q_1, \dots, q_L\}$ where q_1 is the starting state and q_L the final one. Assume that M is matched to a subsequence $X_b^e = \{x_b, \dots, x_e\}$ ($1 \leq b \leq e \leq N$) of X , and that we have an implicit (not modeled) *garbage/filler state* q_G preceding and following M . Thus we implicitly introduce the grammatical constraint that we recognize one keyword occurrence, preceded and followed by non-keyword segments.

Average Observation Posterior

The average observation posterior (AOP) of a model M given a subsequence X_b^e is defined as the average of the log posterior probability along the optimal path, i.e.:

$$\begin{aligned} AOp(M, X, b, e) &= \frac{1}{e-b+1} \min_{Q \in \mathcal{M}} -\log P(Q|X_b^e) \\ &= \frac{1}{e-b+1} \min_{Q \in \mathcal{M}} \{-\log P(q^e|x_e) \\ &\quad - \sum_{n=b}^{e-1} [\log P(q^n|x_n) + \log P(q^{n+1}|q^n)]\} \end{aligned} \quad (1)$$

where $Q = \{q^b, q^{b+1}, \dots, q^e\}$ represents any of the possible paths of length $(e-b+1)$ in M , and q^n the HMM state visited at time n along Q , with $q^n \in \mathcal{Q}$. The local posteriors $P(q_k|x_n)$ can be estimated as output values of a multilayer perceptron (MLP) as used in hybrid HMM/Neural Network systems (Bourlard & Morgan 1994).

For a specific sub-sequence X_b^e , expression (1) can easily be estimated by dynamic programming since the sub-sequence and the associated normalizing factor $(e-b+1)$ are given and constant. However, in the case of keyword spotting, this expression should be estimated for all possible begin/endpoint pairs $\{b, e\}$ (as well as for all possible word models), and we define the matching score of X and M as:

$$AOp(M, X) = AOp(M, X, b_p^*, e_p^*) \quad (2)$$

where the optimal begin/endpoints $\{b_p^*, e_p^*\}$, are the ones for the path Q_p^* yielding the lowest average local log posterior

$$\langle b_p^*, e_p^* \rangle = \operatorname{argmin}_{\{b, e\}} AOp(M, X, b, e) \quad (3)$$

Average Observation Likelihood

The average observation likelihood (AOL) of a subsequence X_b^e given a model M is defined as the average of the log likelihood along the optimal path, i.e.:

$$\begin{aligned} AOL(X, M, b, e) &= \frac{1}{e-b+1} \min_{Q \in \mathcal{M}} -\log P(X_b^e, Q) \\ &= \frac{1}{e-b+1} \min_{Q \in \mathcal{M}} \{-\log P(x_e|q^e) \\ &\quad - \sum_{n=b}^{e-1} [\log P(x_n|q^n) + \log P(q^{n+1}|q^n)]\} \end{aligned} \quad (4)$$

The local likelihoods $P(x_n|q_k)$ can be estimated using gaussian models, etc. (Young 1996). The matching score of X and M according to this criteria is defined as

$$AOL(X, M) = AOL(X, M, b_l^*, e_l^*) \quad (5)$$

where

$$\langle b_l^*, e_l^* \rangle = \operatorname{argmin}_{\{b, e\}} AOL(X, M, b, e) \quad (6)$$

The computation of the AOp and AOL measures is similar, except for the estimation of the local posterior, respectively likelihood. An important factor in deciding which of them should be used is the implementation of the acoustic classifier: gaussians, neural networks, etc. For simplicity in the following we assume that we use posteriors. However, all further discussions apply straightforwardly if posteriors are replaced with likelihoods.

Given the time normalization and the relaxation of begin/endpoints, straightforward DP is no longer optimal and has to be adapted, usually involving more CPU. A new (and simple) solution to this problem is proposed in this paper.

Background

Let us summarize the techniques that compute the same output as the algorithms proposed here or/and that are used as intermediary steps in our techniques.

Filler-based KWS

Although various solutions have been proposed towards the direct optimization of (2) and (5) as, e.g., in (Bridle 1973; Wilpon *et al.* 1989), most of the KWS approaches today

prefer to preserve the simplicity of Viterbi DP by modeling the complete input X (Rohlicek 1995). This is done by explicitly (Rose & Paul 1990) or implicitly (Boite *et al.* 1993) modeling non-keyword segments, using so called filler or garbage models as additional reference models. We assume that non-keyword segments are modeled by extraneous garbage models/states q_G (and grammatical constraints ruling the possible keyword/non-keyword sequences).

In this paper, we will consider only the case of detecting one keyword per utterance at a time. In this case, the keyword spotting problem amounts at matching the whole sequence \bar{X} of length N onto an extended HMM model \bar{M} consisting of the states $\{q_0, q_1, \dots, q_L, q_{L+1}\}$, in which q_0 and q_{L+1} are tied (i.e., share emission probabilities) to the garbage model state q_G . A path of length N in \bar{M} (whose segment contained in M is $Q = \{q^b, \dots, q^e\}$) is denoted

$\bar{Q} = \{\underbrace{q_0, \dots, q_0}_{b-1}, q^b, q^{b+1}, \dots, q^e, \underbrace{q_{L+1}, \dots, q_{L+1}}_{N-e}\}$ with $(b-1)$ garbage states q_0 preceding q^b and $(N-e)$ states q_{L+1} following q^e , and respectively emitting the vector sequences X_1^{b-1} and X_{e+1}^N associated with the non-keyword segments.

Given some estimation ε_t of $-\log P(q_G|x_n)$, the optimal path \bar{Q}_t (and, consequently segmentation b_t and e_t) using the posteriors criteria is then given by:

$$\begin{aligned} \bar{Q}_t &= \underset{\forall \bar{Q} \in \bar{M}}{\operatorname{argmin}} -\log P(\bar{Q}|X) \\ &= \underset{\forall \bar{Q} \in \bar{M}}{\operatorname{argmin}} \{-\log P(Q|X_b^e) \\ &\quad + (b-1) \log P(q_0|q_0) + (N-e) \log P(q_{L+1}|q_{L+1}) \\ &\quad - \sum_{n=1}^{b-1} \log P(q_G|x_n) - \sum_{n=e+1}^N \log P(q_G|x_n)\} \quad (7) \end{aligned}$$

which can be solved by straightforward DP, all paths having the same length (see Algorithm 1). The main problem of filler-based keyword spotting approaches is then to find ways to best estimate $P(q_G|x_n)$ in order to minimize the error introduced by the approximations. In (Boite *et al.* 1993) this value is set to the average of the K best local scores (for some value K). In other approaches it is generated by training explicit filler HMMs. These approaches will usually not lead to the ‘‘optimal’’ solutions given by (2), (5).

Viterbi decoding retrieves the most probable path through a HMM M that generates a certain observation X . Its first step uses a DP algorithm computing the probability of the best path through M given X (i.e., Equation 7). This DP in the first step of Viterbi is isolated in Algorithm 1 (ViterbiScoring), which returns the score of the best path. It uses an array V storing in each $V[i, j]$ the score of the best path of length i in M that starts at the start state and ends in state q_j . For the so called criteria of Maximum A Posteriori (MAP), Algorithm 1 uses local posteriors. For Maximum Likelihood Estimation (MLE), it uses local likelihoods. (Jelinek 1999; Boulard & Morgan 1994)

Sliding Model Method

The sliding model method (Wilpon *et al.* 1989), see Algorithm 2, returns the Average Observation Probability (posterior or likelihood) of the best match of the keyword model

```

function ViterbiScoring( $M, X, Prob$ )
  foreach  $k$  from 2 to  $L, V[1, k] = \infty;$ 
   $V[1, 1] = Prob[1, 1];$ 
  foreach  $i$  from 2 to  $N$  do
     $u = i\%2;$  // index  $i$  is used mod 2 to save space;
    foreach  $j$  from 1 to  $L$  do
       $V[u, j] = Prob[i, j] +$ 
         $\min_{k \in [1..L]} (V[1-u, k] - \log P(q_j|q_k));$ 
  return  $V[N\%2, L];$ 

```

Algorithm 1: Function of whether MLE or MAP is used, $Prob[i, j]$ is $-\log P(x_i|q_j)$ respectively $-\log P(q_j|x_i)$.

```

function SlidingModel( $M, X, Prob$ )
   $score = \infty;$ 
  foreach  $b$  from 1 to  $N-1$  do
    foreach  $k$  from 2 to  $L, V[b\%2, k] = \infty;$ 
     $V[b\%2, 1] = Prob[b, 1];$ 
    foreach  $i$  from  $b+1$  to  $N$  do
       $u = i\%2;$ 
      foreach  $j$  from 1 to  $L$  do
         $V[u, j] = Prob[i, j] +$ 
           $\min_{k \in [1..L]} (V[1-u, k] - \log P(q_j|q_k));$ 
         $score = \min(score, \frac{V[u, L]}{(i-b+1)});$ 
        if  $score$  changed then  $b^* \leftarrow b, e^* \leftarrow i;$ 
  return ( $score, b^*, e^*$ );

```

Algorithm 2: Sliding Model Method.

rior or likelihood) of the best match of the keyword model with a segment of the observation. This is done by computing the score for each segment. It reuses the partial results between segments starting at the same frame. If the basic operation is the computation at Line 2.1, then the cost is $LN(N-1)/2$. $\{b^*, e^*\}$ denote the optimal segmentation.

Segmentation by Filler Re-estimation




Figure 2: Keyword with Filler Model for SFR

In the following we define a method referred to as *Segmentation by Filler Re-estimation (SFR)* with good/fast convergence properties, estimating a value for $-\log P(q_G|x_n)$ (or $-\log P(x_n|q_G)$ when using likelihoods) such that straightforward DP for computing (7) yields exactly the same segmentation (and recognition results) than (3). While the same result is achieved through the sliding model method the algorithm proposed below is more efficient.

```

function SegmFillReest( $M, X, Prob, \varepsilon$ )
  foreach  $k$  from 1 to  $L+1$  do  $V[0, k] = \infty$ ;
   $V[0, 0] = 0$ ;
  while true do
3.1   foreach  $i$  from 1 to  $N$  do
       $u = i \% 2$ ;
       $Prob[i, 0] = Prob[i, L+1] = \varepsilon$ ;
      foreach  $j$  from 0 to  $L+1$  do
3.2          $v = \underset{k \in [0..L+1]}{\operatorname{argmin}} (V[1-u, k] - \log P(q_j|q_k))$ ;
           $b[u, j] = (v == 0) ? i : b[1-u, v]$ ;
           $V[u, j] = Prob[i, j] +$ 
3.3              $(V[1-u, v] - \log P(q_j|q_v))$ ;
3.4      $e[u] = (v \leq L) ? i : e[1-u]$ ;
       $\varepsilon = \frac{V[u, L+1] - \varepsilon N}{e[u] - b[u, L+1]} + \varepsilon$ ;
      if  $\varepsilon$  unchanged then break;
  return  $(\varepsilon, b[u, L+1], e[u])$ ;

```

Algorithm 3: Segmentation by Filler Re-estimation. The ε parameter may take any value. Theoretically the algorithm may be faster if ε is closer to the result. Experimentation shown that it made little difference in practice.

SFR is based on the same HMM structure as the filler based approaches (7). However, rather than looking for explicit (and empirical) estimates of $P(q_G|x_n)$ we aim at mathematically estimating its value (which will be different and adapted to each utterance) such that solving (7) is equivalent to solving (3). Here q_G is a special emitting state whose emission probability density function is not normalized, but is constant (see Figure 2). Also, the transition probabilities from/to garbage states are not normalized (equaling 1). As such, continuing to use the probability notations with (7) would be an abuse of language and we rewrite it using ε instead of $-\log P(q_G|x_n)$, ε having only an algebraic significance. The estimation of ε at step t is denoted ε_t .

$$\overline{Q}_t = \underset{\forall \overline{Q} \in \overline{M}}{\operatorname{argmin}} \{-\log P(\overline{Q}|X_b^e) + (b-1+N-e)\varepsilon_t\} \quad (8)$$

Thus, we perform a converging re-estimation of ε , such that the segmentation resulting of (8) is the same than what would be obtained from (3), and we sometimes called it Iterating Viterbi Decoding (IVD) (Silaghi 1999).

SFR² can be summarized as follows (see Algorithm 3):

1. Start from an initial value³ ε_0 , (e.g., with ε_0 equal to 0 or to a cheap estimation of the score of a “match”).
2. Find \overline{Q}_t (and b_t, e_t) according to (8) (Lines 3.1 to 3.3).
3. Update the estimated value of ε to the average of the local posteriors along the keyword segment of the path Q_t , (Line 3.4) i.e.,:

$$\varepsilon_{t+1} = AOp(M, X, b_t, e_t) = \frac{-1}{(e_t - b_t + 1)} \log P(Q_t|X_{b_t}^{e_t})$$

²Patent (Silaghi 1999)

³In Appendix, it is actually proven that SFR will always converge to the same solution (in more or less cycles, with the worst case upper bound of N loops) independently of this initialization.

4. If $\varepsilon_{t+1} \neq \varepsilon_t$, then increment t and return to Step 2.

Lemma 1 (AOp(M,X) is a fix point) *If $\varepsilon_t = AOp(M, X)$ given by (2), then solving (7) yields $\langle Q^*, b^*, e^* \rangle$ given by (3).*

We note from the definition of AOp that $\forall b, e$,

$$AOp(M, X, b, e) \geq AOp(M, X, b^*, e^*). \quad (9)$$

Lemma 2 (decrease if $\varepsilon_t > AOp(M,X)$) *If $\varepsilon_t > AOp(M, X)$, then the path computed with (8) is \overline{Q}_t with $AOp(M, X, b_t, e_t) < \varepsilon_t$.*

Theorem 1 *The series $\{\langle Q_t, b_t, e_t \rangle\}_t$, defined with the recursion $\varepsilon_{t+1} = AOp(M, X, b_t, e_t)$ where Q_t, b_t, e_t are computed from ε_t with (8), converges to the unique fix point $\langle Q^*, b^*, e^* \rangle$.*

Lemma 3 (shrinking path) *If $\varepsilon_t > \varepsilon_{t+1}$, then either $e_{t+1} - b_{t+1} < e_t - b_t$ or we have reached convergence.*

The proof of these propositions appears in Appendix: each SFR cycle (from the second one) decreases the estimation of ε , and the final path yields the same solution than (3).

Corollary 1.1 *SFR converges in maximum N cycles.*

Proof. From the previous lemma and from Lemma 2 we notice that at each step before convergence (starting with the 2^{nd}), the length of Q_t decreases. The number of steps is therefore roughly upper bounded by N . \square

Another way of proving the convergence of SFR (which yields weaker bounds for termination) consists in proving the equivalence of SFR with a slightly less efficient version where ε is subtracted in each cycle from each element of *Prob*. This new version can then be proven equivalent to an instance of (W.Dinkelbach 1967)’s procedure for a *zero-one fractional programming problem* with real coefficients. Details appear in (Silaghi 2005), as well as application to other criteria.

Since SFR is proven to terminate in at most N cycles, its proven worst cost is $N^2(L+2)$ basic operations, where the basic operation in SFR adds Line 3.2 to the basic operation of the Algorithm 2. While SFR’s theoretical worst case is more than twice worse than the fix (best case) cost of Algorithm 2, SFR always converged in 3 cycles in the experiments described further, leading to an expected cost of $3N(L+2)$, which is much better (N ’s value is between 176 and 1097 in our database). Moreover, absence of a keyword can be decided from the first iteration (see the next section).

Decision by Filler Re-estimation

A typical usage of keyword spotting is to detect the presence of some keyword in an utterance (without caring for the segmentation). If the score of the matching is above some threshold T , then a positive detection is declared.

Lemma 4 *When the estimation of ε is initialized with T , if it increases at the first iteration then the match will be rejected, otherwise it will be accepted.*

Proof. From (9) it follows that whenever ε_0 has a smaller value than the fix point, the first re-estimation will transform it in a value bigger than the fix point, i.e., $\varepsilon_1 > \varepsilon_0$. If ε_0 has

a larger value than the fix point, the first re-estimation will yield a smaller value, $\varepsilon_1 < \varepsilon_0$ (Lemma 2). Therefore the direction of change in the first re-estimation of ε tells if the initialization value was smaller or bigger than the fix point.

If we only want to know how T compares to the fix point (which is the score of the match), it follows that we can learn this in one cycle by setting ε_0 to T . \square

```

function DecisionFillReest( $M, X, Prob, T$ )
  foreach  $k$  from 1 to  $L+1$  do  $V[0, k] = \infty$ ;
   $V[0, 0] = 0$ ;
  foreach  $i$  from 1 to  $N$  do
     $u = i \% 2$ ;
     $Prob[i, 0] = Prob[i, L+1] = T$ ;
    foreach  $j$  from 0 to  $L+1$  do
       $V[u, j] = Prob[i, j] +$ 
         $\min_{k \in [0..L+1]} (V[1-u, k] - \log P(q_j | q_k))$ ;
  return  $V[u, L] > TN$ ;

```

Algorithm 4: Decision by Filler Re-estimation.

The obtained algorithm is called *Decision by Filler Re-estimation (DFR)* (see Algorithm 4) and roughly consists of a cycle of SFR. The basic operation (defined as for the sliding model method) is performed only $N(L+2)$ times.

Experimental results

Our algorithms and the sliding model method were compared on the BREF database (Lamel, Gauvain, & Eskénazi 1991), a continuous, read speech microphone database. 242 utterances (with a 2,300 word lexicon), were used for testing. These keywords were represented by simple hybrid HMM/ANN models based on context-independent phones.

On this database, the average speedup of SFR (13.867.986 DP updates/keyword) vs. sliding model (1.241.899.326 updates/keyword) is 89.5 times. SFR gives the same results as the sliding model method (confirming the soundness of the implementation). For computing the segmentation, 3 iterations were needed in each experiment. While it may be believed that the fast convergence was due to lucky choices of databases and keywords, the same fast convergence was observed for all keywords in the database and for also for a very different application to natural language parsing of another version of our algorithm (Rosenknop & Silaghi 2001).

DFR is SFR which stops after the first (out of 3) iterations. DFR is obviously 3 times faster than SFR and yields the same result. For building a ROC curve it is preferable to use SFR since DFR would have to be run again for each computed point of the curve. However, for production systems (working at a predefined point) DFR is preferable.

The use of likelihoods versus posteriors is dictated by the technology used for the probabilistic estimation, since ANNs return posteriors and most other learning techniques compute efficiently only likelihoods. It is more rigorous to use posteriors than likelihoods, since the likelihoods-based criteria is an approximation of the posterior-based criteria.

Certain results suggest that posteriors perform also experimentally better than likelihoods (Bouclard & Morgan 1994). However, experimental methods have not yet been devised to separate the influence of the difficult to evaluate quality of the training of different estimators.

We do not provide here the comparison between recognition with AOP scoring versus other scorings, since these comparisons have been thoroughly studied elsewhere (Wilpon *et al.* 1990; Silaghi & Bouclard 2000). In a summary, some other scores work better than AOP with certain keywords and databases, and vice-versa. AOP's advantage is that it is not a black-box but a rigorous measure and its behavior improves understanding of keyword spotting. Other criteria that we proposed are the double normalization and the real fitting (Silaghi 1999).

Conclusions

We have thus proposed new methods for keyword spotting (KWS), working with both local likelihoods and local posterior probabilities, and optimizing rigorously defined matching scores (the average observation probability) between keyword models and keyword segments. The fastest previous algorithm with this property is Sliding Model requiring exactly $N^2L/2$ basic operations, a prohibitive cost.

A first proposed algorithm, *Segmentation by Filler Re-estimation (SFR)*, computes the same output like Sliding Model but much faster: $3N(L+2)$ measured in two very different experimented applications, KWS and NLP, even if its theoretical worst case cost is $O(N^2L)$. A second proposed algorithm, called *Decision by Filler Re-estimation (DFR)* returns the same matching decision like SFR and sliding model methods, but does not compute the segmentation (which is not required in many applications). However its cost is fix, $N(L+2)$ of the same basic operations like in sliding model method. Proof is given in Appendix. Typical applications are search in audio databases and recognition of spoken passwords, but due to its use of the general HMM framework, the proposed algorithms have been shown to have applications in computer vision, natural language processing, and bioinformatics (Silaghi 1999; Rosenknop & Silaghi 2001).

Acknowledgements

We thank Hervé Bouclard for the trained ANNs used in experiments, and Giulia Bernardis for help in setting up the BREF database.

Appendix

Lemma 1 (AOP(M,X) is a fix point) *If $\varepsilon_t = AOp(M, X)$ given by (2), then solving (7) yields $\langle Q^*, b^*, e^* \rangle$ given by (3).*

Proof. Let us note $AOp(M, X)$ with w . We see from definitions that $-\log P(Q^*|X) = N * w$. A suboptimal path Q_b^e would have yielded $-\log P(Q_b^e|X) = N * w + (e - b + 1) * (w' - w) > N * w$ since $w' > w$ from the definition of AOp . Here we have noted by w' the value of the average log probability in the suboptimal path Q_b^e , i.e.:

$$w' = AOp(M, X, b, e) = \frac{1}{e - b + 1} - \log P(Q_b^e|X_b^e) > w$$

Since DP yields the optimum, it chooses $\overline{Q^*}$. \square

Lemma 2 (decrease if $\varepsilon_t > \mathbf{AOp}(M, X)$) *If $\varepsilon_t > \mathbf{AOp}(M, X)$, then the path computed with (8) is $\overline{Q_t}$ with $\mathbf{AOp}(M, X, b_t, e_t) < \varepsilon_t$.*

Proof. Let us denote $\mathbf{AOp}(M, X, b_t, e_t)$ with w' and $\mathbf{AOp}(M, X)$ with w .

$-\log P(\overline{Q_t}|X) = (N-1-e_t+b_t)*\varepsilon_t + (e_t-b_t+1)*w'$
(i.e., contribution of non-keyword segments '+' contribution of keyword segment). Since the path $\overline{Q^*}$ was not preferred:

$$-\log P(\overline{Q_t}|X) \leq -\log P(\overline{Q^*}|X)$$

This can be written as: $(N-1-e_t+b_t)*\varepsilon_t + (e_t-b_t+1)*w' \leq (N-1-e^*+b^*)*\varepsilon_t + (e^*-b^*+1)*w$.

After reordering we obtain: $(e_t-b_t+1)*w' \leq (N-1-e^*+b^*)*\varepsilon_t + (e^*-b^*+1)*w - (N-1-e_t+b_t)*\varepsilon_t$, and by regrouping the right term:

$$(e_t-b_t+1)*w' \leq (e^*-b^*+1)*w + ((e_t-b_t)-(e^*-b^*))*\varepsilon_t \quad (10)$$

Since $e^*-b^*+1 > 0$ and from the hypothesis that $w < \varepsilon_t$:

$$(e^*-b^*+1)*w + ((e_t-b_t)-(e^*-b^*))*\varepsilon_t < (e_t-b_t+1)*\varepsilon_t \quad (11)$$

From the inequalities (10) and (11):

$$(e_t-b_t+1)*w' < (e_t-b_t+1)*\varepsilon_t$$

$$e_t-b_t+1 > 0 \Rightarrow w' < \varepsilon_t.$$

If $\overline{Q^*}$ was preferred, then again $w' = w < \varepsilon_t$. \square

Theorem 1 *The series $\{\langle Q_t, b_t, e_t \rangle\}_t$, defined with the recursion $\varepsilon_{t+1} = \mathbf{AOp}(M, X, b_t, e_t)$ where Q_t, b_t, e_t are computed from ε_t with (8), converges to the unique fix point $\langle Q^*, b^*, e^* \rangle$.*

Proof. From inequality (9), after the 2nd step we have $\varepsilon \geq \mathbf{AOp}(M, X)$. From Lemma 2, each further cycle of the algorithm decreases ε .

From Lemma 1 and Lemma 2, results that the convergence appears at $\mathbf{AOp}(M, X, b_t, e_t) = \mathbf{AOp}(M, X)$. Since the number of possible paths is finite, it converges. \square

Lemma 3 (shrinking path) *If $\varepsilon_t > \varepsilon_{t+1}$, then either $e_{t+1} - b_{t+1} < e_t - b_t$ or we have reached convergence.*

Proof. Let us denote $w_t = \mathbf{AOp}(M, X, b_t, e_t)$ and $w_{t+1} = \mathbf{AOp}(M, X, b_{t+1}, e_{t+1})$. Since $\overline{Q_t}$ was preferred to $\overline{Q_{t+1}}$ for $\varepsilon = \varepsilon_t$, it means that $\varepsilon_t(b_t - e_t + N - 1) + w_t(e_t - b_t + 1) \leq \varepsilon_t(b_{t+1} - e_{t+1} + N - 1) + w_{t+1}(e_{t+1} - b_{t+1} + 1)$.
 $\varepsilon_t((e_{t+1} - b_{t+1}) - (e_t - b_t)) + (e_t - b_t + 1)w_t - (e_{t+1} - b_{t+1} + 1)w_{t+1} \leq 0$ (12)

If for $\varepsilon = \varepsilon_{t+1}$ ($\varepsilon_{t+1} < \varepsilon_t$), we would obtain $\overline{Q_{\varepsilon_{t+1}}}$ with $e_{t+1} - b_{t+1} \geq e_t - b_t$, $(e_{t+1} - b_{t+1}) - (e_t - b_t) \geq 0$, then by adding $(\varepsilon_{t+1} - \varepsilon_t)((e_{t+1} - b_{t+1}) - (e_t - b_t)) \leq 0$ to inequality (12) we obtain:

$$\varepsilon_{t+1}((e_{t+1} - b_{t+1}) - (e_t - b_t)) + (e_t - b_t + 1)w_t - (e_{t+1} - b_{t+1} + 1)w_{t+1} \leq 0$$

showing that $\overline{Q_{\varepsilon_t}}$ will still be preferred to $\overline{Q_{\varepsilon_{t+1}}}$. If $\overline{Q_{\varepsilon_t}} \neq \overline{Q_{\varepsilon_{t+1}}}$ then this is a contradiction and we can infer that $e_{t+1} - b_{t+1} < e_t - b_t$. Otherwise we have $\overline{Q_{\varepsilon_t}} = \overline{Q_{\varepsilon_{t+1}}}$ and we reach convergence with $e_{t+1} - b_{t+1} = e_t - b_t$. \square

References

- Bellman, R. 1952. On the theory of dynamic programming. *Proc. of the National Academy of Sciences* 38:716–719.
- Benayed, Y.; Fohr, D.; Haton, J.; and Chollet, G. 2003. Confidence measures for keyword spotting using support vector machines. In *ICASSP*.
- Boite, J.-M.; Boulard, H.; d'Hoore, B.; and Haesen, M. 1993. A new approach towards keyword spotting. In *EUROSPEECH*, 1273–1276.
- Boulard, H., and Morgan, N. 1994. *Connectionist Speech Recognition - A Hybrid Approach*. Kluwer.
- Bridle, J. 1973. An efficient elastic-template method for detecting given words in running speech. In *Brit. Acoust. Soc. Meeting*, 1–4.
- Jelinek, F. 1999. *Statistical Methods for Speech Recognition*. MIT Press.
- Junkawitsch, J.; Neubauer, L.; Höge, H.; and Ruske, G. 1996. A new keyword spotting algorithm with pre-calculated optimal thresholds. In *ICSLP*, 2067–2070.
- Lamel, L.-F.; Gauvain, J.-L.; and Eskénazi, E. 1991. Bref, a large vocabulary spoken corpus for french. In *Proceedings of EuroSpeech'91*.
- Rohlicek, J. 1995. Word spotting. In Ramachandran, R., and Mammone, R., eds., *Modern Methods of Speech Processing*. Kluwer. 123–157.
- Rose, R., and Paul, D. 1990. A hidden markov model based keyword recognition system. In *ICASSP'90*, 129–132.
- Rosenknop, A., and Silaghi, M.-C. 2001. Algorithme de décodage de treillis selon le critère de coût moyen pour la reconnaissance de la parole. In *TALN'01*.
- Russell, S., and Norvig, P. 2002. *Artificial Intelligence, A Modern Approach*. Prentice Hall, 2nd edition.
- Silaghi, M.-C., and Boulard, H. 2000. Iterative posterior-based keyword spotting without filler models. In *Proceedings of the IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing*.
- Silaghi, M.-C. 1999. Speech recogn. and sign. anal. by search of subsequences with maximal confidence measure. PCT patent WO 00/51107.
- Silaghi, M. C. 2005. Iterative segmentation in keyword spotting: relation to fractional programming. Technical Report CS-2005-11, FIT.
- W.Dinkelbach. 1967. On nonlinear fractional programming. *Management Science* 18(7):492–498.
- Wilpon, J.; Rabiner, L.; C.-H., L.; and Goldman, E. 1989. Application of hidden markov models of keywords in unconstrained speech. In *Proc. of ICASSP'89*, 254–257.
- Wilpon, J.; Rabiner, L.; Lee, C.; and Goldman, E. 1990. Automatic recognition of keywords in unconstrained speech using hidden markov models. *IEEE Trans. on ASSP* 38(11):1870–1878.
- Young, S. 1996. A review of large-vocabulary continuous-speech recognition. *IEEE Signal Proc. Magazine* 45–57.