UNSUPERVISED SPOKEN TERM DETECTION WITH SPOKEN QUERIES BY MULTI-LEVEL ACOUSTIC PATTERNS WITH VARYING MODEL GRANULARITY

Cheng-Tao Chung^{#1}, *Chun-an Chan*^{*2}, *and Lin-shan Lee*^{#*3}

Graduate Institute of Electrical Engineering, National Taiwan University[#] Graduate Institute of Communication Engineering, National Taiwan University^{*}

b97901182@gmail.com¹, chunanchan@gmail.com², lslee@gate.sinica.edu.tw³

ABSTRACT

This paper presents a new approach for unsupervised Spoken Term Detection with spoken queries using multiple sets of acoustic patterns automatically discovered from the target corpus. The different pattern HMM configurations(number of states per model, number of distinct models, number of Gaussians per state)form a three-dimensional model granularity space. Different sets of acoustic patterns automatically discovered on different points properly distributed over this three-dimensional space are complementary to one another, thus can jointly capture the characteristics of the spoken terms. By representing the spoken content and spoken query as sequences of acoustic patterns, a series of approaches for matching the pattern index sequences while considering the signal variations are developed. In this way, not only the on-line computation load can be reduced, but the signal distributions caused by different speakers and acoustic conditions can be reasonably taken care of. The results indicate that this approach significantly outperformed the unsupervised feature-based DTW baseline by 16.16% in mean average precision on the TIMIT corpus.

Index Terms— zero resource speech recognition, unsupervised learning, dynamic time warping, hidden Markov models, spoken term detection

1. INTRODUCTION

The fast growing quantity of video and audio content over the Internet implies a very high demand for efficient and accurate approaches to search through the spoken contents. Spoken term detection (STD) usually refers to the task of finding all occurrences of the text query term from a large spoken archive [1]. Most STD approaches were based on automatic speech recognition (ASR), transforming speech into words or subwords for token matching [1][2][3][4], with performance relying heavily on the ASR accuracy [5]. This implies annotated training corpora properly matched to the spoken content are necessary. When the input query is spoken, it becomes possible to directly match the spoken content with the spoken query without conventional ASR. In this way, the difficulties in conventional ASR such as recognition errors and need for annotated training data may be bypassed, which is especially attractive for languages with very limited annotated data [6][7] or spoken content with unknown languages. This leads to recent efforts in unsupervised STD with spoken queries without using annotated data during training [8] [9], which is also the focus of this work. Hereafter we assume all queries are in spoken, and no annotated speech data is available.

Prevailing approaches to the task considered here rely on dynamic time warping (DTW) to directly match the spoken query to the spoken documents. However, a major limitation of DTW-based approaches is that the DTW distances are easily affected by speaker mismatch and varying acoustic conditions. Many related works focused on feature representations and distance measures within the DTW framework that are more robust to speaker and acoustic condition variations[10], including using the posteriorgrams from a universal Gaussian mixture model [11], and the acoustic segment models [12][13]. The other limitation for DTW-based approaches is that the computation load for the matching process is linear to the number of frames to be searched through. Substantial efforts were devoted to reducing this computation load, such as segment-based DTW [9], lowerbound estimation for DTW [14][15], and a locality sensitive hashing technique for indexing speech frames [16].

In recent years substantial effort has been made for unsupervised model based discovery of acoustic patterns from corpora without manual annotation[17][18][19][20][21][22]. In this paper, we propose a new approach for unsupervised STD with spoken queries using multilevel acoustic patterns automatically discovered from the target corpus with varying model granularity discovery from the corpus of interest. The different pattern HMM configurations(number of states per model, number of distinct models, number of Gaussians per state)form a threedimensional model granularity space. Different sets of acoustic patterns automatically discovered on different points properly distributed over this three-dimensional space are complementary to on another, thus can jointly capture the characteristics of the spoken terms. By converting the spoken content and query into parallel sequences of acoustic patterns with different model granularity, token matching can be performed with pattern indices representing highly varying signals. Very encouraging results were obtained the preliminary experiments.

2. ACOUSTIC PATTERNS WITH VARYING MODEL GRANULARITY

2.1. Pattern Discovery for a Given Model Configuration

Given an unlabelled speech corpus, it is not difficult for unsupervised discovery of the desired acoustic patterns from the corpus for a chosen hyperparameter set ψ that determines the topology of the HMMs [19][20][21]. This can be achieved by first finding an initial label ω_0 based on a set of assumed patterns for all observations in the corpus χ as in (1)[17]. Then in each iteration t the HMM parameter set θ_t^{ψ} can be trained with the label ω_{t-1} obtained in the previous iteration as in (2), and the new label ω_t can be obtained by free-pattern decoding with the obtained parameter set θ_t^{ψ} as in (3).

$$\omega_0 = \text{initialization}(\chi), \tag{1}$$

$$\theta_t^{\psi} = \arg\max_{\theta^{\psi}} P(\chi | \theta^{\psi}, \omega_{t-1}), \tag{2}$$

$$\omega_t = \arg \max P(\chi | \theta_t^{\psi}, \omega). \tag{3}$$

The training process can be repeated with enough number of iterations until the difference between ω_{t-1} and ω_t becomes insignificant. This gives a converged set of acoustic pattern HMMs which we denote as Θ^{ψ} .

2.2. Model Granularity Space

The above process can be performed with many different HMM configurations, each characterized by three hyperparameters: the number of states m in each acoustic pattern HMM, the total number of acoustic patterns n during initialization, and the number of Gaussians l in each HMM state, $\psi = (m, n, l)$. The transcription of a speech signal decoded with these acoustic pattern HMMs may be considered as a temporal segmentation of the signal, so the HMM length (or number of states in each HMM) m represents the temporal granularity. The set of all acoustic pattern HMMs may be considered as a segmentation of the phonetic space, so the total number n of acoustic pattern HMMs represents the phonetic granularity. The different Gaussians in each state then jointly model the distributions of the signals in the acoustic feature space represented by MFCCs, so the number of Gaussians l in each state represents the acoustic granularity. This gives a three-dimensional representation of the acoustic pattern configurations in terms of temporal, phonetic and acoustic granularities as in Fig. 1. Any point in the three-dimensional space in Fig. 1 corresponds to an acoustic pattern configuration.

Although the hyperparameters $\psi = (m, n, l)$ are difficult to determine for a given corpus of unknown language and unknown linguistic characteristics, it is possible to have many different sets of acoustic patterns with hyperparameters and HMMs { $\Theta^{\psi_k}, \psi_k = (m_k, n_k, l_k), k = 1, 2, ...K$ } independently learned in parallel, considered as on different levels. Because different model granularities (m, n, l) give different characteristics to the acoustic patterns as mentioned above, with enough number of pattern sets and the model granularities (m, n, l) properly distributed in the three-dimensional space, this multi-level set of acoustic pattern may jointly represent the behavior of the signal for the given corpus. There can be a variety of applications for these patterns, but here in the work below we use these patterns in spoken term detection, assuming the characteristics of the spoken terms can be captured by different sets of these patterns.



Fig. 1: The model granularity space for acoustic pattern configurations.

3. SPOKEN TERM DETECTION AND SEARCH METHODS

3.1. Off-line Processing

All the spoken documents in the archive are first off-line decoded into sequences of acoustic patterns using each level of acoustic pattern HMM set $\Theta^{(m_k,n_k,l_k)}$. Let $\{p_r, r = 1, 2, 3, ..., n_k\}$ denote the n_k acoustic

patterns in the set Θ^{ψ_k} . We further construct a similarity matrix S of size $n_k \times n_k$ for every ψ_k , for which the component S(i, j) is the similarity between any two pattern HMMs p_i and p_j in the set Θ^{ψ_k} . Two similarity matrices used for this work are in (4).

$$S(i,j) = \begin{cases} \delta(i,j), & \text{for hard similarity, or} \quad (4a) \\ \exp(-\mathrm{KL}(i,j)/\beta, & \text{for soft similarity.} \quad (4b) \end{cases}$$

The matrix in (4a) is simply the identity matrix. The KL-divergence KL(i, j) between two pattern HMMs in (4b) is defined as the KL-divergence between the states based on the variational approximation [23] summed over the states. To transform the KL divergence into a similarity measure between 0 and 1, a negative exponential was applied [24] with a scaling factor β . When β is small, similarity between distinct patterns in (4b) approaches zero, so (4b) becomes similar to (4a). β can be determined with a held out data set, but here we simply set it to 100 times the number of states m_k .

3.2. On-line Matching Matrix Construction

In the on-line phase, we perform the following for each entered spoken query q and each document d in the archive. Assume for the pattern set Θ^{ψ_k} a document d is decoded into a sequence of D acoustic patterns with indices $(d_1, d_2, ..., d_D)$ and the query q into a sequence of Q patterns with indices $(q_1, ..., q_Q)$. We thus construct a matching matrix W of size $D \times Q$ for every document-query pair, in which each entry (i, j)is the similarity between acoustic patterns with indices d_i and q_j as in (5a).

$$W(i,j) = \begin{cases} S(d_i, q_j), & \text{for 1-best sequences, or} & (5a) \\ P_i^T SP_j, & \text{for N-best sequences.} & (5b) \end{cases}$$

Alternatively, we can also match the N-best sequences of documents to N-best sequences of queries as depicted in Fig (2) and (5b). We extend each pattern position *i* in a sequence into a posteriorgram vector P_i of size $n_k \times 1$ by accumulating the duration for each of the n_k patterns within the pattern boundaries of the best transcription, across the N-best transcriptions considered, uniformly weighted and normalized(P_i for *d* and P_j for *q*). When only the best transcription is considered, (5b) reduces to (5a). The price paid here is the increased computation time by a factor of $O(n_k^2)$ over that in (5a). One can also choose to use the posteriorgram only on the query, which increases the load by $O(n_k)$.



Fig. 2: Construction of pattern posteriorgram from N-best list

3.3. On-line Matching Policy

There can be two methods to calculate the relevance score between document d and query q. In the sub-sequence matching(SUB) method in (6a), we sum the elements in the matrix W along the diagonal direction, generating the accumulated similarities for all sub-sequences starting at all pattern positions in d as shown in Fig (3)(a). The maximum is selected to represent the relevance between document d and query q on the pattern set Θ^{ψ_k} as in (6a).

$$R(d,q) = \begin{cases} \max_{i:\text{index}} \sum_{j=1}^{Q} W(i+j,j), & \text{for SUB, or (6a)} \\ \\ \max_{u:\text{ path}} \sum_{j=1}^{|u|} W(u_d(j), u_q(j)), & \text{for DTW. (6b)} \end{cases}$$

In order to alleviate the problem of insertion/deletion, we can also perform dynamic time warping (DTW) on the matrix W as in (6b) and Fig 3(b), refered to as the pattern-based DTW here. This is an extended version of (6a), except now the elements W are accumulated along any allowed DTW path in the matrix W. The summation in (6b) is over a single DTW path, while the maximization in (6b) is over all DTW paths. Although DTW takes longer time, performing DTW in the matrix W on-line is significantly faster than the conventional frame-based DTW, because most of the calculation was performed offline when evaluating (4). Since table lookup is constant time operation, asymptotically speaking, the computational load online is reduced by a factor of $O(FT^2) = O(Fm_k^2)$, where F is the feature dimension in the framebased DTW and T is the duration of the acoustic patterns in frames which scales linearly with number of states m_k in HMMs.



Fig. 3: The matching matrix W

3.4. Overall Relevance Score

In each of (4)(5)(6) there are two options, leading to a total 8 search methods. We thus use three binary digits to specify these methods in reporting experimental results below: $\gamma = (\text{Soft,Nbest,DTW})$, i.e. Soft=1 for soft similarity in (4b) and 0 for (4a); Nbest=1 for N-best sequence in (5b) and 0 for (5a); DTW=1 for DTW in (6b) and 0 for (6a). These search methods, $\{\gamma_s, s = 1, ..., 8\}$ give different relevance scores for each pattern set Θ^{ψ_k} , $R^{(\psi_k, \gamma_s)}(d, q)$ as in (6). The overall relevance score $\bar{R}(d, q)$ between d and q is then simply the weighted sum of (6) over all K different sets of acoustic patterns Θ^{ψ_k} and the 8 search methods with weights $\lambda(\psi_k, \gamma_s)$ as in (7). Below, we simply set the weights $\lambda(\psi_k, \gamma_s)$ to either 0 or 1.

$$\bar{R}(d,q) = \sum_{k=1}^{K} \sum_{s=1}^{8} \lambda(\psi_k, \gamma_s) R^{(\psi_k, \gamma_s)}(d,q).$$
(7)

Note that the acoustic pattern sets $\{\Theta^{\psi_k}, k = 1, 2, ..., K\}$ for different model granularities are complementary to one another. By adding the scores obtained via different pattern sets as in (7) the signal characteristics can be better captured. In addition, the limitation caused by temporal granularity (model length m_k) may be alleviated to some degree by the pattern-based DTW in (6b), the limitation caused by phonetic granularity (number of patterns n_k) may be alleviated to some degree by the Nbest sequences in (5b), and the limitation caused by acoustic granularity (number of Gaussian l_k) may be taken care to some degree by the soft similarity in (4b). Therefore, the choices of ψ_k and γ_s are correlated for good scores of $\overline{R}(d,q)$ in (7). Also, the process above, including constructing the KL-divergence matrices and decoding the spoken documents into acoustic patterns off-line, and generating the pattern sequence for the query on-line can all be performed in a highly parallel manner, so the computational load can be scalable with the computational resources available.

4. EXPERIMENTS

The proposed approach was tested in preliminary experiments performed on the TIMIT corpus. 20 sets of acoustic patterns with number of states m=3, 5, 7, 9, 11 and number of HMMs n=50, 100, 200, 300 were first generated with l=1 Gaussian per state on the TIMIT training set. Then we increased the number of Gaussians per state by 1 and perform (2) and (3) until the sets converge. We repeated the process until we had l=1, 2, 3, 4 for all the 20 pattern sets, obtaining K=80 pattern HMM sets in the end. With the 8 search methods mentioned in section 3, this gave a total of 640 scores for every query-document pair in (7). The TIMIT training set was also taken as the spoken archive from which we wish to detect the spoken terms.

The query set consisted of 32 spoken words randomly selected from the TIMIT testing set. An instance of every query word was randomly selected from the testing set, and used as the spoken query to search for other instances in the training set. Note that although choices of $\lambda(\psi_k, \gamma_s)$ can be based on the optimization with respect to an evaluation metric [25][26], the preliminary experiments in this work were to verify the feasibility of the proposed frameworks instead. The baseline we compared to was frame-based DTW on MFCC sequences. The same acoustic features were used for training the pattern HMMs. In principle, the framework should generalize to other features as well. For spoken term detection the performance measures we used were the mean average precision(MAP), precision at 5 and 10(P@5 and P@10). All three measures gave very similar trends. Below we only report results for MAP due to space limitation.

4.1. Feature Selection and Achievable Performance

In this experiment, our goal was to learn the 640 weights $\lambda(\psi, \gamma)$ in (7) to be either 1 or 0 in order to optimize the MAP. We randomly split the query set into 2 disjoint sets A and B, each containing 16 queries. We use set B as the development set for learning these weights and focus our discussion on the performance on set A. Starting with every $\lambda(\psi, \gamma)=0$, we greedily select the next (ψ, γ) that would yield the best MAP and set it to 1. This process was repeated until 20 pairs of (ψ, γ) were selected. The results are shown in pink in Fig. 4, in which the oracle results by learning on set A itself are also shown in cyan. The baseline of framebased DTW was 10.16% for set A. It was low probably because most queries unexpectedly had less than 10 relevant documents in the training set, and the various dialects of TIMIT made it even more difficult for simply comparing the feature sequences. We can clearly see with only 20 out of 640 scores selected based on set B, the MAP reaches 25.88%, significantly higher than the baseline. The oracle results learning on set A itself reached as high as 35.60%, which implied that if a more elaborate learning algorithm was applied, there was still much room for improvement.

The parameter sets (ψ, γ) for the 20 scores selected based on set B are also printed on Fig. 4. Some observations can be made here. In all the 20 cases Soft=1, so Soft=1 was certainly better. For the selection of N-best sequences or DTW there were no clear trends. However, some correlation between search methods and pattern configurations can be observed. For example, it can be found that Nbest=1 were preferred when l=1 or l=2, probably because too few number of Gaussians limited the accuracies in the 1-best sequence. Also, DTW=1 was preferred very often when m=5 or 3, probably because for shorter patterns it made better sense to merge more than one patterns into a longer pattern, which is actually what the pattern-based DTW did. These results seem to imply

the need to learn from a development set, which is not always available. This will be further discussed in the next section.



Fig. 4: MAP performance on query set A with top 20 scores selected from a development set.

4.2. Performance Analysis without a Development Set

Here we consider the case without a development set. We first consider the different search method γ by summing all the 80 scores for all combinations of m,n,l, but not γ as shown in Part (A) of Fig. (5). Several conclusions may be drawn: (a) Soft similarity brought massive improvement (Soft=1 > Soft=0 for all combinations of Nbest and DTW), (b) N-best Sequences brought only a negligible improvement (Nbest=1 \sim Nbest=0 for all combinations of Soft and DTW), (c) DTW degraded the performance in general (DTW=1 < DTW=0 for all combinations of Soft and Nbest). Conclusion (a) is consistent with the conclusion drawn from the top 20 scores in Fig 4. Since generating a soft similarity metric is also the only part of the search that could be conducted off-line, it is certainly attractive. Conclusion (b) and (c) may be a little surprising, since intuitively N-best sequences and DTW can help and quite several of the top 20 scores in Fig 4 had Nbest=1 or DTW=1. As discussed previously with Fig 4, Nbest and DTW could be helpful for specific pattern configurations ψ but not necessarily all. When such specific configurations cannot be properly chosen with a development set, these improvements could be diluted when averaged with all possible configurations. Because the different pattern sets carried complementary information, jointly considering the 80 1-best sequences obtained from the 80 pattern sets itself can be viewed as considering a single large lattice best representing the utterance with all time warping and N-best information included. This is probably why N-best and DTW didn't help here. The MAP obtained by summing all 640 scores without any selection was 20.62% as shown in Part(A) of Fig (5), which was also significantly higher than the baseline.

Therefore below we focus our discussion on the selection of pattern configurations of ψ assuming $\gamma = (1,0,0)$ without using Nbest or DTW. The 80 scores for different ψ forms a 3 dimensional space over (m, n, l). We summed the relevance score over 2 of the dimensions and plotted the performance on the remaining dimension. Summing over (m, n), (m, l), (n, l), we get Parts (B)(C)(D) of Fig. (5) respectively. The average of the MAP values for Parts (B)(C)(D) of Fig. (5) are also listed for comparison between the dimensions. Several additional conclusions may be drawn: (d) the performance is better for larger l as shown in Part(B) of Fig. 5, (e) Model combination on the (m, n) plane was the most effective (the average MAP of Part(B) was much higher than those in Part (C)(D) of Fig. 5), (f) several optimal values seem to exist for m and n (the maximum occurs for m=3 in Part (D) and n=100 in Part (C)). We highly suspect that these optimal points for m, n as mentioned in Conclusion (f) are inert characteristics of the underlying language that should stay approximately the same for a different corpus of the same

language. Conclusion (d) will probably hold until the phenomenon of over-fitting begins to happen, since the number of Gaussians is limited by the training corpus size. Conclusion (e) implies when blending score with respect to (m, n), a good policy may be to have m and n as diverse as possible, if there is no information regarding the selection of m, and n. From conclusion (f), m=3 seemed close to the average phoneme duration of TIMIT, while n=100 seemed close to the number of phonemes with some context dependency considered. This may imply selection of (m, n) is a language dependent characteristic although we do not have strong evidence yet to back up this claim. This would be useful if verified to be true, especially for under-resourced languages, since in that case the learned weights could be similarly useful for different corpora on the same language.



Fig. 5: MAP performance on query set A without a development set

We further plot the MAP of $R^{(\psi,\gamma)}(d,q)$ for γ =(1,0,0) and l=3 for sets A and B in Fig 6, the best performing γ and l in Parts (A)(B) of Fig 5. The 20 points for (m, n) were interpolated with a 2D spline function to show the smoothed MAP distributions over the plane. As can be seen, the score distributions look similar to a good degree even for completely different query sets in Fig 6(a) and (b). When we simply selected the 20 scores for set A here (m=3,5,7,9,11; n=50,100,200,300; l=3, γ =(1,0,0)), the MAP is 26.32%, slightly higher than 25.88% achieved above by the 20 scores selected greedily from a development set.



Fig. 6: MAP for γ =(1,0,0) and *l*=3 over the (*m*, *n*) plane for query sets A and B.

5. CONCLUSION

In this work, we present a new approach for unsupervised spoken term detection using multi-level acoustic patterns discovered from the target corpus. The different pattern sets with different model configurations are complementary, thus can jointly capture the information for the spoken terms. Significantly better performance than frame-based DTW on TIMIT corpus as obtained.

6. REFERENCES

- David RH Miller, Michael Kleber, Chia-Lin Kao, Owen Kimball, Thomas Colthurst, Stephen A Lowe, Richard M Schwartz, and Herbert Gish, "Rapid and accurate spoken term detection.," in *INTERSPEECH*, 2007, pp. 314–317.
- [2] Jonathan Mamou, Bhuvana Ramabhadran, and Olivier Siohan, "Vocabulary independent spoken term detection," in *Proceedings* of the 30th annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 2007, pp. 615–622.
- [3] Roy G Wallace, Robert J Vogt, and Sridha Sridharan, "A phonetic search approach to the 2006 nist spoken term detection evaluation," 2007.
- [4] Yi-cheng Pan and Lin-shan Lee, "Performance analysis for lattice-based speech indexing approaches using words and subword units," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 18, no. 6, pp. 1562–1574, 2010.
- [5] Murat Saraclar and Richard Sproat, "Lattice-based search for spoken utterance retrieval," *Urbana*, vol. 51, pp. 61801, 2004.
- [6] Lou Boves, Rolf Carlson, Erhard W Hinrichs, David House, Steven Krauwer, Lothar Lemnitzer, Martti Vainio, and Peter Wittenburg, "Resources for speech research: present and future infrastructure needs.," in *INTERSPEECH*. Citeseer, 2009, pp. 1803–1806.
- [7] Arun Kumar, Nitendra Rajput, Dipanjan Chakraborty, Sheetal K Agarwal, and Amit A Nanavati, "Wwtw: the world wide telecom web," in *Proceedings of the 2007 workshop on Networked systems* for developing regions. ACM, 2007, p. 7.
- [8] Florian Metze, Nitendra Rajput, Xavier Anguera, Marelie Davel, Guillaume Gravier, Charl Van Heerden, Gautam V Mantena, Armando Muscariello, Kishore Prahallad, Igor Szoke, et al., "The spoken web search task at mediaeval 2011," in Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on. IEEE, 2012, pp. 5165–5168.
- [9] Chun-An Chan and Lin-Shan Lee, "Unsupervised hidden markov modeling of spoken queries for spoken term detection without speech recognition.," in *INTERSPEECH*, 2011, pp. 2141–2144.
- [10] Michael A Carlin, Samuel Thomas, Aren Jansen, and Hynek Hermansky, "Rapid evaluation of speech representations for spoken term discovery.," in *INTERSPEECH*, 2011, pp. 821–824.
- [11] Yaodong Zhang and James R Glass, "Unsupervised spoken keyword spotting via segmental dtw on gaussian posteriorgrams," in Automatic Speech Recognition & Understanding, 2009. ASRU 2009. IEEE Workshop on. IEEE, 2009, pp. 398–403.
- [12] Marijn Huijbregts, Mitchell McLaren, and David van Leeuwen, "Unsupervised acoustic sub-word unit detection for query-byexample spoken term detection," in Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on. IEEE, 2011, pp. 4436–4439.
- [13] Haipeng Wang, Cheung-Chi Leung, Tan Lee, Bin Ma, and Haizhou Li, "An acoustic segment modeling approach to query-by-example spoken term detection," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on.* IEEE, 2012, pp. 5157–5160.

- [14] Yaodong Zhang and James R Glass, "A piecewise aggregate approximation lower-bound estimate for posteriorgram-based dynamic time warping," in *INTERSPEECH*, 2011, pp. 1909–1912.
- [15] Yaodong Zhang, Kiarash Adl, and James Glass, "Fast spoken query detection using lower-bound dynamic time warping on graphical processing units," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on.* IEEE, 2012, pp. 5173–5176.
- [16] Aren Jansen and Benjamin Van Durme, "Indexing raw acoustic features for scalable zero resource search.," in *INTERSPEECH*, 2012.
- [17] Cheng-Tao Chung, Chun-an Chan, and Lin-shan Lee, "Unsupervised discovery of linguistic structure including two-level acoustic patterns using three cascaded stages of iterative optimization," in Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on. IEEE, 2013, pp. 8081–8085.
- [18] Aren Jansen, Kenneth Church, and Hynek Hermansky, "Towards spoken term discovery at scale with zero resources.," in *INTER-SPEECH*, 2010, pp. 1676–1679.
- [19] Aren Jansen and Kenneth Church, "Towards unsupervised training of speaker independent acoustic models.," in *INTERSPEECH*, 2011, pp. 1693–1692.
- [20] Herbert Gish, Man-hung Siu, Arthur Chan, and William Belfield, "Unsupervised training of an hmm-based speech recognizer for topic classification.," in *INTERSPEECH*, 2009, pp. 1935–1938.
- [21] Mathias Creutz and Krista Lagus, "Unsupervised models for morpheme segmentation and morphology learning," ACM Transactions on Speech and Language Processing (TSLP), vol. 4, no. 1, pp. 3, 2007.
- [22] Hung-yi Lee, Yun-Chiao Li, Cheng-Tao Chung, and Lin-shan Lee, "Enhancing query expansion for semantic retrieval of spoken content with automatically discovered acoustic patterns," in Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on. IEEE, 2013, pp. 8297–8301.
- [23] John R Hershey and Peder A Olsen, "Approximating the kullback leibler divergence between gaussian mixture models," in Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on. IEEE, 2007, vol. 4, pp. IV–317.
- [24] Marcin Marszałek and Cordelia Schmid, "Constructing category hierarchies for visual recognition," in *Computer Vision–ECCV* 2008, pp. 479–491. Springer, 2008.
- [25] Filip Radlinski and Thorsten Joachims, "Query chains: learning to rank from implicit feedback," in *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. ACM, 2005, pp. 239–248.
- [26] Tie-Yan Liu, "Learning to rank for information retrieval," Foundations and Trends in Information Retrieval, vol. 3, no. 3, pp. 225– 331, 2009.