Unsupervised Learning of Tone and Pitch Accent

Gina-Anne Levow

Department of Computer Science
University of Chicago, IL 60637 USA
levow@cs.uchicago.edu

Abstract

Recognition of tone and intonation is essential for speech recognition and language understanding. However, most approaches to this recognition task have relied upon extensive collections of manually tagged data obtained at substantial time and financial cost. In this paper, we explore unsupervised clustering approaches to recognize pitch accent in English and tones in Mandarin Chinese. In unsupervised Mandarin tone clustering experiments, we achieve 57-87% accuracy on materials ranging from broadcast news to clean lab speech. For English pitch accent in broadcast news materials, results reach 78%. These results indicate that the intrinsic structure of tone and pitch accent acoustics can be exploited to reduce the need for costly labeled training data for tone learning and recognition.

1. Introduction

Tone and intonation play a crucial role across many languages. However, the use and structure of tone varies widely, ranging from lexical tone which determines word identity to pitch accent signalling information status. Here we consider the recognition of lexical tones in Mandarin Chinese syllables and pitch accent in English.

Although intonation is an integral part of language and is requisite for understanding, recognition of tone and pitch accent remains a challenging problem. The majority of current approaches to tone recognition in Mandarin and other East Asian tone languages integrate tone identification with the general task of speech recognition within a Hidden Markov Model framework. In some cases tone recognition is done only implicitly when a word or syllable is constrained jointly by the segmental acoustics and a higher level language model and the word identity determines tone identity. Other strategies build explicit and distinct models for the syllable final region, the vowel and optionally a final nasal, for each tone.

Recent research has demonstrated the importance of contextual and coarticulatory influences on the surface realization of tones.[1] The overall shape of the tone or accent can be substantially modified by the local effects of adjacent tone and intonational elements. Furthermore, broad scale phenomena such as topic and phrase structure can affect pitch height, and pitch shape may be variably affected by the presence of boundary tones. These findings have led to explicit modeling of tonal context within the HMM framework. In addition to earlier approaches that employed phrase structure, several recent approaches to tone recognition in East Asian languages [2, 3] have incorporated elements of local and broad range contextual influence on tone. Many of these techniques create explicit context-dependent models of the phone, tone, or accent for each context in which they appear, either using the tone sequence for left or right context or using a simplified high-low contrast, as is natural for integration in a Hidden Markov Model speech recognition framework. In pitch accent recognition, recent work by [4] has integrated pitch accent and boundary tone recognition with speech recognition using prosodically conditioned models within an HMM framework, improving both speech and prosodic recognition.

Since these approaches are integrated with HMM speech recognition models, standard HMM training procedures which rely upon large labeled training sets are used for tone recognition as well. Other tone and pitch accent recognition approaches using other classification frameworks such as support vector machines [5] and decision trees with boosting and bagging [6] have relied upon large labeled training sets - thousands of instances - for classifier learning. This labelled training data is costly to construct, both in terms of time and money, with estimates for some intonation annotation tasks reaching tens of times real-time. This annotation bottleneck as well as a theoretical interest in the learning of tone motivates the use of unsupervised or semi-supervised approaches to tone recognition whereby the reliance on this often scarce resource can be reduced.

Little research has been done in the application of unsupervised techniques for tone and pitch accent recognition. Some preliminary work by [7] employs self-organizing maps and measures of f0 velocity for tone learning. In this paper we explore the use of spectral and standard k-means clustering for unsupervised acquisition of tone. We find that in clean read speech, unsupervised techniques can identify the underlying Mandarin tone categories with high accuracy, while even on noisier broadcast news speech, Mandarin tones can be recognized well above chance levels, with English pitch accent recognition at 96% of the levels achieved with fully supervised Support Vector Machine (SVM) classifiers.

The remainder of paper is organized as follows. Section 2 describes the data sets on which English pitch accent and Mandarin tone learning are performed and the feature extraction process. Section 3 describes the unsupervised techniques employed and Section 4 the experiments and results. Section 5 presents conclusions and future work.

2. Data Sets

We consider two corpora: one in English for pitch accent recognition and two in Mandarin for tone recognition. We introduce each briefly below.

2.1. English Corpus

We employ a subset of the Boston Radio News Corpus [8], read by female speaker F2B, comprising 40 minutes of news material. The corpus includes pitch accent, phrase and boundary

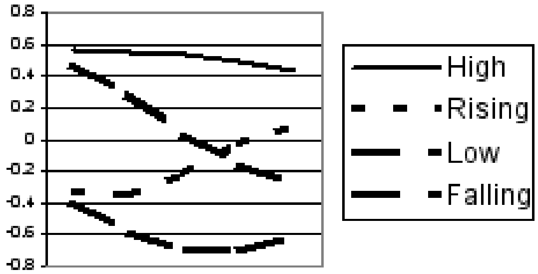

Figure 1: Contours for canonical Mandarin tones

tone annotation in the ToBI framework [9] aligned with manual transcription and syllabification of the materials. Following earlier research [10, 6], we collapse the ToBI pitch accent labels to four classes: unaccented, high, low, and downstepped high for experimentation.

2.2. Mandarin Chinese Tone Data

Mandarin Chinese is a language with lexical tone in which each syllable carries a tone and the meaning of the syllable is jointly determined by the tone and segmental information. Mandarin Chinese has four canonical lexical tones, typically described as follows: 1) high level, 2) mid-rising, 3) low falling-rising, and 4) high falling.[1] The canonical pitch contours for these tones appear in Figure 1. We employ data from two distinct sources in the experiments reported here.

2.2.1. Read Speech

The first data set is very clean speech data drawn from a collection of read speech collected under laboratory conditions by [11]. In these materials, speakers read a set of short sentences where syllable tone and position of focus were varied to assess the effects of focus position on tone realization. Focus here corresponds to narrow focus, where speakers were asked to emphasize a particular word or syllable. Tones on focussed syllables were found to conform closely to the canonical shapes described above, and in previous supervised experiments using a linear support vector machine classifier trained on focused syllables, accuracy approached 99%. For these materials, pitch tracks were manually aligned to the syllable and automatically smoothed and time-normalized by the original researcher, resulting in 20 pitch values for each syllable.

2.2.2. Broadcast News Speech

The second data set is drawn from the Voice of America Mandarin broadcast news, distributed by the Linguistic Data Consortium[2], as part of the Topic Detection and Tracking (TDT-2) evaluation. Using the corresponding anchor scripts, automatically word-segmented, as gold standard transcription, audio from the news stories was force-aligned to the text transcripts. The forced alignment employed the language porting functionality of the University of Colorado Sonic speech recognizer [12]. A mapping from the transcriptions to English phone sequences supported by Sonic was created using a Chinese character-pinyin pronunciation dictionary and a manually constructed mapping from pinyin sequences to the closest corresponding English phone sequences.[3]

2.3. Acoustic Features

We employ a common representation for both tone and pitch accent recognition. In prior supervised experiments using support vector machines, this representation achieved competitive recognition levels for tone and pitch accent recognition. The representation is described below.

Using Praat's [13] ”To pitch” and ”To intensity” functions and the alignments generated above, we extract acoustic features for the prosodic region of interest. This region corresponds to the “final” region of each syllable in Chinese, including the vowel and any following nasal, and to the syllable nucleus in English.[4] For all pitch and intensity features in both datasets, we compute per-speaker z-score normalized log-scaled values. We extract pitch values from five points evenly spaced across valid pitch tracked points in the syllable. We also compute mean pitch across the syllable. Recent phonetic research [14] has identified significant effects of carryover coarticulation from preceding adjacent syllable tones. To minimize these effects consistent with the pitch target approximation model [14], we compute slope features based on the second half of this final region, where this model predicts that the underlying pitch height and slope targets of the syllable will be most accurately approached. We further log-scale and normalize slope values to compensate for greater speeds of pitch fall than pitch rise[15].

3. Unsupervised Learning

The bottleneck of time and monetary cost associated with manual annotation has generated significant interest in the development of techniques for machine learning and classification that reduce the amount of annotated data required for training. Likewise, learning from unlabeled data aligns with the perspective of language acquisition, as child learners must identify these linguistic categories without explicit instruction by observation of natural language interaction. Of particular interest are techniques in unsupervised learning where the structure of unlabeled examples may be exploited. Here we consider unsupervised techniques with no labeled training data.

A wide variety of unsupervised clustering techniques have been proposed. In addition to classic clustering techniques such as k-means, recent work has shown good results for many forms of spectral clustering including those by [16, 17, 18]. In the unsupervised experiments reported here, we employ asymmetric k-lines clustering by [18] using code available at the authors' site, as our primary unsupervised learning approach. Asymmetric clustering is distinguished from other techniques by the construction and use of context-dependent kernel radii. Rather than assuming that all clusters are uniform and spherical, this approach enhances clustering effectiveness when clusters may not be spherical and may vary in size and shape. We will see that this flexibility yields a good match to the structure of Mandarin tone data where both shape and size of clusters vary across tones. In additional contrastive experiments reported below,

[1] For the experiments in this paper, we exclude the neutral tone, which appears on unstressed syllables.

[2] http://www.ldc.upenn.edu

[3] All tone transformations due to third tone sandhi are applied to create the label set.

[4] We restrict our experiments to syllables with at least 50 ms of tracked pitch in this final region.

we also compare k-means clustering, symmetric k-lines clustering [18], and Laplacian Eigenmaps [17] with k-lines clustering. The spectral techniques all perform spectral decomposition on some representation of the affinity or adjacency graph.

We contrast results under unsupervised learning with most common class assignment and previous results employing fully supervised approaches, such as SVMs.

4. Unsupervised Clustering Experiments

We executed four sets of experiments in unsupervised clustering using the [18] asymmetric clustering algorithm.

4.1. Experiment Configuration

In these experiments, we chose increasingly difficult and natural test materials. In the first experiment with the cleanest data, we used only focused syllables from the read Mandarin speech dataset. In the second, we included both in-focus (focused) and pre-focus syllables from the read Mandarin speech dataset.[5] In the third and fourth experiments, we chose subsets of broadcast news report data, from the Voice of America (VOA) in Mandarin and Boston University Radio News corpus in English.

In all experiments on Mandarin data, we performed clustering on a balanced sampling set of tones, with 100 instances from each class[6], yielding a baseline for assignment of a single class to all instances of 25%. We then employed a two-stage repeated clustering process, creating 2 or 3 clusters at each stage.

For experiments on English data, we extracted a set of 1000 instances, sampling pitch accent types according to their frequency in the collection. We performed a single clustering phase with 2 to 16 clusters, reporting results at different numbers of clusters.

For evaluation, we report accuracy based on assigning the most frequent class label in each cluster to all members of the cluster.

4.2. Experimental Results

We find that in all cases, accuracy based on the asymmetric clustering is significantly better than most common class assignment and in some cases approaches 96% of labelled classification accuracy. Unsurprisingly, the best results, in absolute terms, are achieved on the clean focused syllables, reaching 87% accuracy. For combined in-focus and pre-focus syllables, this rate drops to 77%. These rates contrast with 99-93% accuracies in supervised classification using linear SVM classifiers with several thousand labelled training examples.

On broadcast news audio, accuracy for Mandarin reaches 57%, still much better than the 25% level, though below a 72% accuracy achieved using supervised linear SVMs with 600 labeled training examples. Interestingly, for English pitch accent recognition, accuracy reaches 78.4%, approximately 96% of the 81.3% accuracy achieved with SVMs on a comparable data representation.

4.3. Contrastive Experiments

We further contrast the use of different unsupervised learners, comparing the three spectral techniques and k-means with Euclidean distance. All contrasts are presented for English pitch accent classification, ranging over different numbers of clusters, with the best parameter setting of neighborhood size. The results are illustrated in Figure 2. Results for k-means and the asymmetric clustering technique are presented for the clean focal Mandarin speech under the standard two stage clustering.

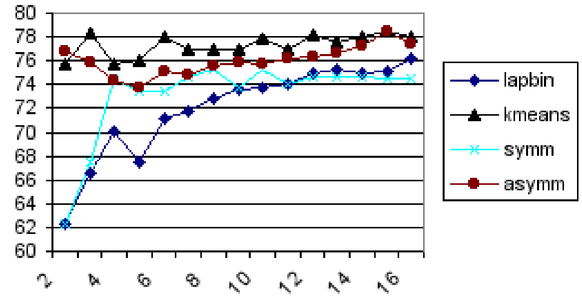

Figure 2: Differences for alternative unsupervised learners across numbers of clusters.

The asymmetric k-lines clustering approach consistently outperforms the corresponding symmetric clustering learner, as well as Laplacian Eigenmaps with binary weights for pitch accent classification. Somewhat surprisingly, k-means clustering outperforms all of the other approaches when producing 3-14 clusters. Accuracy for the optimal choice of clusters and parameters is comparable for asymmetric k-lines clustering and k-means, and somewhat better than all other techniques considered. The careful feature selection process for tone and pitch accent modeling may reduce the difference between the spectral and k-means approaches. In contrast, for the four tone classification task in Mandarin using two stage clustering, the best clustering using asymmetric k-lines strongly outperforms k-means, at 87% and 74.75% accuracy respectively.

4.4. Discussion

An examination of both the clusters formed and the structure of the data provides insight into the effectiveness of this process. Figure 3 displays 2 dimensions of the Mandarin four-tone data from the focused read speech, where normalized pitch mean is on the x-axis and slope is on the y-axis. The separation of classes and their structure is clear. One observes that rising tone (tone 2) lies above the x-axis, while high-level (tone 1) lies along the x-axis. Low (tone 3) and falling (tone 4) tones lie mostly below the x-axis as they generally have falling slope. Low tone (3) appears to the left of falling tone (4) in the figure, corresponding to differences in mean pitch.

In clustering experiments, an initial 2- or 3-way split separates falling from rising or level tones based on pitch slope. The second stage of clustering splits either by slope (tones 1,2, some 3) or by pitch height (tones 3,4). These clusters capture the natural structure of the data where tones are characterized by pitch height and slope targets.

5. Conclusion & Future Work

We have demonstrated the effectiveness of unsupervised techniques for recognition of Mandarin Chinese syllable tones and English pitch accents using acoustic features alone to capture pitch target height and slope. Although outperformed by fully supervised classification techniques using much larger samples

[5] Post-focus syllables typically have decrease pitch height and range, resulting in particularly poor recognition accuracy. We chose not to concentrate on this specific tone modeling problem here.

[6] Sample sizes were bounded to support rapid repeated experimentation and for consistency with the relatively small VOA data set.

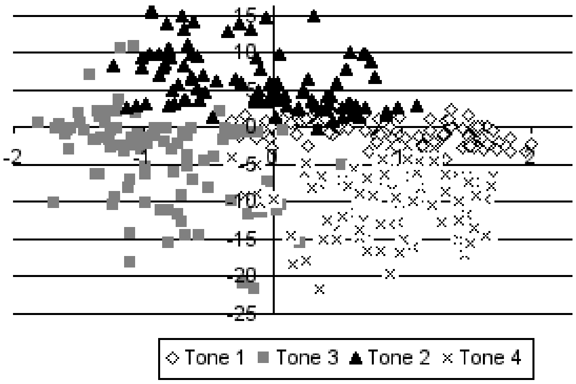

Figure 3: Scatterplot of pitch height vs pitch slope. Open Diamond: High tone (1), Filled black traingle: Rising tone (2), Filled grey square: Low tone (3), X: Falling tone (4)

	Unsup.	Supervised
Mandarin Tone		
Lab, In-focus	87%	99%
Lab, Pre & In-focus	77%	93%
Broacast News	57%	72%
English Pitch Accent		
Broadcast News, 4 class	78%	81.3%

Table 1: Summary of experimental results

of labelled training data, these unsupervised techniques perform well above most common class assignment, in the best cases approaching supervised levels. Table 1 presents a summary of experimental results. Unsupervised techniques achieve accuracies of 87% on the cleanest read speech, reaching 57% on data from a standard Mandarin broadcast news corpus, and over 78% on pitch accent classification for English broadcast news.

Future work will consider a broader range of tone and intonation classification, including the richer tone set of Cantonese as well as Bantu family tone languages, where annotated data truly is very rare. We also hope to integrate a richer contextual representation of tone and intonation consistent with phonetic theory within this unsupervised framework.

6. Acknowledgements

We would like to thank Yi Xu for granting access to the read speech data, Vikas Sindhwani, Mikhail Belkin, and Partha Niyogi for their implementation of Laplacian Eigenmaps, and Igor Fischer and J. Poland for their implementation of asymmetric clustering.

7. References

[1] Yi Xu, "Contextual tonal variations in Mandarin," *Journal of Phonetics*, vol. 25, pp. 62–83, 1997.

[2] C. Wang and S. Seneff, "Improved tone recognition by normalizing for coarticulation and intonation effects," in *Proceedings of 6th International Conference on Spoken Language Processing*, 2000.

[3] J. L. Zhou, Ye Tian, Yu Shi, Chao Huang, and Eric Chang, "Tone articulation modeling for Mandarin spontaneous speech recognition," in *Proceedings of ICASSP 2004*, 2004.

[4] M. Hasegawa-Johnson, Jennifer Cole, Chilin Shih abd Ken Chen, Aaron Cohen, Sandra Chavarria, Heejin Kim, Taejin Yoon, Sarah Borys, and Jeung-Yoon Choi, "Speech recognition models of the interdependence among syntax, prosody, and segmental acoustics," in *HLT/NAACL-2004*, 2004.

[5] Nuttakorn Thubthong and Boonserm Kijsirikul, "Support vector machines for Thai phoneme recognition," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 9, no. 6, pp. 803–813, 2001.

[6] Xuejing Sun, "Pitch accent prediction using ensemble machine learning," in *Proceedings of ICSLP-2002*, 2002.

[7] Bruno Gauthier, Rushen Shi, Yi Xu, and Robert Proulx, "Neural-network simulation of tonal categorization based on f0 velocity profiles," *Journal of the Acoustical Society of America*, vol. 117, Pt. 2, pp. 2430, 2005.

[8] M. Ostendorf, P. J. Price, and S. Shattuck-Hufnagel, "The Boston University radio news corpus," Tech. Rep. ECS-95-001, Boston University, 1995.

[9] K. Silverman, M. Beckman, J. Pitrelli, M. Ostendorf, C. Wightman, P. Price, J. Pierrehumbert, and J. Hirschberg, "ToBI: A standard for labelling English prosody," in *Proceedings of ICSLP*, 1992, pp. 867–870.

[10] M. Ostendorf and K. Ross, "A multi-level model for recognition of intonation labels," in *Computing Prosody*, Y. Sagisaka, N. Campbell, and N. Higuchi, Eds., pp. 291–308. 1997.

[11] Y. Xu, "Effects of tone and focus on the formation and alignment of f0 contours - evidence from Mandarin," *Journal of Phonetics*, vol. 27, 1999.

[12] B. Pellom, W. Ward, J. Hansen, K. Hacioglu, J. Zhang, X. Yu, and S. Pradhan, "University of Colorado dialog systems for travel and navigation," 2001.

[13] P. Boersma, "Praat, a system for doing phonetics by computer," *Glot International*, vol. 5, no. 9–10, pp. 341–345, 2001.

[14] C.X. Xu, Y. Xu, and L.-S. Luo, "A pitch target approximation model for f0 contours in Mandarin," in *Proceedings of the 14th International Congress of Phonetic Sciences*, 1999, pp. 2359–2362.

[15] Yi Xu and X. Sun, "Maximum speed of pitch change and how it may relate to speech," *Journal of the Acoustical Society of America*, vol. 111, 2002.

[16] Jianbo Shi and Jitendra Malik, "Normalized cuts and image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, 2000.

[17] Mikhail Belkin and Partha Niyogi, "Laplacian eigenmaps and spectral techniques for embedding and clustering," in *Proceeding of NIPS'02*, 2002.

[18] I. Fischer and J. Poland, "New methods for spectral clustering," Tech. Rep. ISDIA-12-04, IDSIA, 2004.