

Supporting Teacher Assessment in Chinese Language Learning Using Textual and Tonal Features

Ashvini Varatharaj^{1(\boxtimes)}, Anthony F. Botelho¹, Xiwen Lu², and Neil T. Heffernan¹

 Worcester Polytechnic Institute, WPI, Worcester, MA 01609, USA {avaratharaj,abotelho,nth}@wpi.edu
 ² Brandeis University, 415 South Street, Waltham, MA 02453, USA

xiwenlu@brandeis.edu

Abstract. Assessment in the context of foreign language learning can be difficult and time-consuming for instructors. Distinctive from other domains, language learning often requires teachers to assess each student's ability to speak the language, making this process even more timeconsuming in large classrooms which are particularly common in postsecondary settings; considering that language instructors often assess students through assignments requiring recorded audio, a lack of tools to support such teachers makes providing individual feedback even more challenging. In this work, we seek to explore the development of tools to automatically assess audio responses within a college-level Chinese language-learning course. We build a model designed to grade student audio assignments with the purpose of incorporating such a model into tools focused on helping both teachers and students in real classrooms. Building upon our prior work which explored features extracted from audio, the goal of this work is to explore additional features derived from tone and speech recognition models to help assess students on two outcomes commonly observed in language learning classes: fluency and accuracy of speech. In addition to the exploration of features, this work explores the application of Siamese deep learning models for this assessment task. We find that models utilizing tonal features exhibit higher predictive performance of student fluency while text-based features derived from speech recognition models exhibit higher predictive performance of student accuracy of speech.

Keywords: Audio processing \cdot Natural language processing \cdot Language learning

1 Introduction

When learning a new language, it is important to be able to assess proficiency in skills pertaining to both reading and speaking; this is true for instructors but also for students to understand where improvement is needed. The ability to read requires an ability to identify the characters and words correctly, while successful speech requires correct pronunciation and, in many languages, correctness of tone. For these reasons, read-aloud tasks, where students are required to speak while following a given reading prompt, are considered an integral part of any Standardized language testing system for the syntactic, semantic, and phonological understanding that is required to perform the task well [1–3].

This aspect of learning a second language is particularly important in the context of learning Mandarin Chinese. Given that Chinese (Mandarin Chinese) is a tonal language, the way the words are pronounced could change the entire meaning of the sentence, highlighting the importance of assessing student speech (through recordings or otherwise) an important aspect of understanding a student's proficiency in the language.

While a notable amount of research has been conducted in the area of automating grading of read-aloud tasks by a number of organizations (cf. the Educational Testing Service's (ETS) and Test of English as a Foreign Language (TOEFL)), the majority of assessment of student reading and speech is not taking place in standardized testing centers, but rather in classrooms. It is therefore here, in these classroom settings, that better tools are needed to support both teachers and students in these assessment tasks. In the current classroom paradigm, it is not unreasonable to estimate that the teacher takes hours to listen to the recorded audios and grade them; a class of 20 students providing audio recordings of just 3 min each, for example, requires an hour for the teacher to listen, and this does not include the necessary time to provide feedback to students.

This work observes student data collected from a college-level Chinese language learning course. We use data collected in form of recorded student audio from reading assignments with the goal of developing models to better support teachers and students in assessing proficiency in both fluency, a measure of the coherence of speech, and accuracy, a measure of lexical correctness. We present a set of analyses to compare models built with audio, textual, and tonal features derived from openly available speech-to-text tools to predict both fluency and accuracy grades provided by a Chinese language instructor.

2 Related Work

There has been little work done on developing tools to support the automatic assessment of speaking skills in a classroom setting, particularly in foreign language courses. However, a number of approaches have been applied in studying audio assessment in non-classroom contexts. Pronunciation instruction through computer-assistance tools has received attention by several of the standardized language testing organizations including ETS, SRI, and Pearson [4] in the context of such standardized tests; much of this work is similarly focused on English as a second language learners.

In developing models that are able to assess fluency and accuracy of speech from audio, it is vital for such models to utilize the right set of representative features. Previous work conducted in the area of Chinese language learning, of which this work is building upon, explored a number of commonly-applied features of audio including spectral features, audio frequency statistics, as well as others [5]. Other works have previously explored the similarity and differences between aspects of speech. In [6], phonetic distance, based on the similarity of sounds within words, has been used to analyse non-native English learners. Work has also been done on analysing different phonetic distances used to analyse speech recognition of vocabulary and grammar optimization [7].

Many approaches having been explored in such works, starting from Hidden Markov Models [8] to more recently applying deep learning methods [6] to predict scores assessing speaking skills. Others have utilized speech recognition techniques for audio assessment. There have been a number of prior works that have focused on the grading of read-aloud and writing tasks using Automatic Speech Recognition using interactive Spoken Dialogue Systems (SDS) [9], phonetic segmentation [10,11], as well as classification [12] and regression tree (CART) based methodologies [8,13]. Some speech recognition systems have used language-specific attributes, such as tonal features, to improve their model performances [14,15]. Since tones are an important component of pronunciation in Chinese language learning, we also consider the use of tonal features in this work for the task of predicting teacher-provided scores.

As there has been seemingly more research conducted in the area of natural language processing, such an approach as to convert the spoken audio to text is plausibly useful in understanding the weak points of the speaker. Recent works have combined automatic speech recognition with natural language processing to build grading models for English Language [16]. In applications of natural language processing, the use of pre-trained word embeddings has become more common due to the large corpuses of data on which they were trained. Pretrained models of word2vec [17] and Global Vectors for Word Representation (GloVe) [18], for example, have been widely cited in applications of natural language processing. By training on large datasets, these embeddings are believed to capture meaningful syntactic and semantic relationships between words through meaning. Similar to these methods, FastText [19] is a library created by Facebook's AI Research lab which provides pre-trained word embeddings for Chinese language.

Similarly, openly available speech-generation tools may be useful in assessing student speech. For example, Google has supplied an interface to allow for textto-speech generation that we will utilize later in this work. Though this tool is not equivalent to a native Chinese-speaking person, such a tool may be useful in helping to compare speech across students. The use of this method will be discussed further in Sect. 4.5.

3 Dataset

The data set used in this work was obtained from an undergraduate Chinese Language class taught to non-native speakers. The data was collected from multiple classes with the same instructor. The data is comprised of assignments requiring students to submit an audio recording of them reading aloud a predetermined prompt as well as answering open-ended questions. For this work, we focus only on the read-aloud part of the assignment, observing the audio recordings in conjunction with the provided text prompt of the assignment.

For the read-aloud part of the assignment, students were presented with 4 tasks which were meant to be read out loud and recorded by the student and submitted through the course's learning management system. Each task consists of one or two sentences about general topics. The instructor downloaded such audio files, listened to each, and assessed students based on two separate grades pertaining to fluency and accuracy of the spoken text. Each of these grades are represented as a continuous-valued measure between 0-5; decimal values are allowed such that a grade of 2.5 is the equivalent of a grade of 50% on a particular outcome measure. This dataset contains 304 audio files from 128 distinct students over four distinct sentence read-aloud tasks. Each audio is taken as a separate data point, so each student has one to four audio files. Each sample includes one of the tasks read by the student along with the intended text of the reading prompt.

4 Feature Extraction

4.1 Pre-processing

The audio files submitted by students were of varying formats including mp3, m4a. The ffmpeg [20] python library is used to convert these audio inputs into a raw .wav format required by the speech models.

While prior work [5] explored much of this process of processing and extracting features from audio, the current work intends to expand upon this prior work by additionally introducing textual and tonal features into our grading model. The next two sections detail the feature extraction process for audio, tonal. and textual features using the processed dataset.

4.2 Audio Features

Audio feature extraction is required to obtain components of the audio signal which can be used to represent acoustic characteristics in a way a model can understand. The audio files are converted to vectors which can capture the various properties of the audio data recorded. The audio features were extracted using an openly-available python library [21] that breaks each recording into 50 millisecond clips, offsetting each clip by 25 milliseconds of the start of the preceding clip (creating a sliding window to generate aggregated temporal features using the observed frequencies of the audio wave). A total of 34 audio-based features, including Mel-frequency cepstral coefficients (MFCC), Chroma features, and Energy related features, were generated as in previous work [5].

4.3 Text Features

From the audio data, we also generate a character-representation of the interpreted audio file using openly-available speech recognition tools. The goal of this feature extraction step is to use speech recognition to transcribe the words spoken by each student to text that can be compared to the corresponding reading prompt using natural language processing techniques; the intuition here is that the closeness of what the Google speech-to-text model is able to interpret to the actual prompt should be an indication of how well the given text was spoken. Since building speech recognition models is not the goal of this paper, we used an off-the-shelf module for this task. Specifically, the SpeechRecognition library [22] in python provides a coding interface to the Google Web Speech API and also supports several languages including Mandarin Chinese. The API is built using deep-learning models and allows text transcription in real-time, promoting its usage for deployment in classroom settings. While, to the authors' knowledge, there is no detailed documentation describing the precise training procedure for Google's speech recognition model, it is presumably a deep learning model trained on a sizeable dataset; it is this later aspect, the presumably large number of training samples, that we believe may prove some benefit to our application. Given that we have a relatively small dataset, the use of pre-trained models such as those supplied openly through Google, may be able to provide additional predictive power to models utilizing such features.

We segment the Google-transcribed text into character-level components and then convert them into numeric vector representations for use in the models. We use Facebook's FastText [19] library for this embedding task (c.f. Section 2). Each character or word is represented in the form of a 300 dimensional numeric vector space representation. The embedding process results in a character level representation, but what is needed is a representation of the entire sentence. As such, once the embeddings are applied, all characters are concatenated together to form a large vector representing the entire sentence.

The Google Speech to text API is reported to exhibit a mean Word Error Rate (WER) of 9% [23]. To analyse the performance of Google's Speech to Text API on our dataset of student recordings, we randomly selected 15 student audio files from our dataset across all the four reading tasks and then transcribed them using the open tool. We then created a survey that was answered by the Teacher and 2 Teaching Assistants of the observed Chinese language class. The survey first required the participants to listen to each audio and then asked each to rate the accuracy of the corresponding transcribed text on a 10-point integer scale. The Intraclass Correlation Coefficient (ICC) was used to measure the strength of inter-rater agreement, finding a correlation of 0.8 (c.f. ICC(2,k) [24]). While this small study illustrates that the Google API exhibits some degree of error, we argue that it is reliable enough to be used for comparison in this work.

4.4 Tonal Features

Chinese is a tonal language. The same syllable can be pronounced with different tones which, in turn, changes the meaning of the content. To aid in our goal of predicting the teacher-supplied scores of fluency and accuracy, we decided to explore the observance of tonal features in our models. In the Mandarin Chinese language, there are four main tones. These tones represent changes of inflection (i.e. rising, falling, or leveling) when pronouncing each syllable of a word or phrase. When asking Chinese Language teachers what are some of the features they look for while assessing student speech, tonal accuracy was one of the important characteristics identified.

To extract the tones from the student's audio, we use the ToneNet [25] model which was trained on the Syllable Corpus of Standard Chinese Dataset (SCSC). The SCSC dataset consists of 1,275 monosyllabic Chinese characters, which are composed of 15 pronunciations of young men, totaling 19,125 example pronunciations of about 0.5 to 1 second in duration. The model uses a mel spectrogram (image respresentation of an audio in the mel scale) of each of these samples to train the model. The model uses a convolutional neural network and multi-layer perceptron to classify Chinese syllables in the form on images into one of the four tones. This model is reported to have an accuracy of 99.16% and f1-score of 99.11% [25]. To use the ToneNet on our student audio data, we first break the student audio into 1 second audio clips and convert them into mel spectrograms. We then feed these generated mel spectrograms to the ToneNet model to predict the tone present in each clip. The sequence of predicted tones is then used as features in our fluency and accuracy prediction models.

4.5 True Audio: Google Text to Speech API

As a final source of features for comparison in this work, we believed it may be useful to compare each student audio to that of an accepted "correct" pronunciation; however, no such recordings were present in our data, nor are they common to have in classroom settings for a given read-aloud prompt. Given that we have audio data from students, and the text of each corresponding prompt, we wanted to utilise Google's text-to-speech to produce a "true audio" - how Google would read the given sentence. While not quite ground truth, given that it is trained on large datasets, we believe it could help our models learn certain differentiating characteristics of student speech by providing a common point of comparison. The features extracted from the true audio is particularly useful in training Siamese networks, as described in the next section, by providing a reasonable audio recording with which to compare each student response.

5 Models

In developing models to assess students based on the measures of accuracy and fluency, we compare three models of varying complexities and architectures (and one baseline model) using different feature sets described in the previous sections. Our baseline model consists of assigning the mean of the scores as the predicted value. We use a 5-fold cross validation for all model training (Fig. 1).



Fig. 1. The figure shows the steps involved in transforming audio to sequence of tones and feeding into the Siamese network

Aside from the baseline model, the first and second models explored in this work are the same as applied in previous research in developing models for assessing student accuracy and fluency in Chinese language learning [5]. These models consist of a decision tree (Using the CART algorithm [26]) and a Long Short Term Memory (LSTM) recurrent neural network. While previous work explored the use of audio features, labeled in this work as "PyAudio" features after the library used to generate them, this work is able to compare these additional textual and tonal feature sets. Similar to deep learning models, the decision tree model is able to learn non-linear relationships in the data, but also can be restricted in its complexity to avoid potential problems of overfitting. Conversely, the LSTM is able to learn temporal relationships from time series data as in the audio recordings observed in this work. As in our prior research, a small amount of hyperparameter tuning was conducted on a subset of the data.

In addition to the three sets of features described, a fourth feature set, a cosine similarity measure, was explored in the decision tree model. This was calculated by taking the cosine similarity between the embedded student responses and the embedded reading prompts. This feature set was included as an alternative approach to the features described in Sect. 4.5 for use in the Siamese network described in the next section.

5.1 Siamese Network

The last type of model explored in this paper is a Siamese neural network. Siamese networks are able to learn representations and relationships in data by comparing similar examples. For instance, we have the audio of the student as well as the Google-generated "true audio" that can be compared to learn features that may be useful in identifying how differences correlate with assigned fluency and accuracy scores. In this regard, the generated audio does not need to be correct to be useful; they can help in understanding how the student audio recordings differ from each other and how these differences relate to scores.

The network is comprised of two identical sub networks that share the same weights while working on two different inputs in tandem (e.g. the network observes the student audio data at the same time that it observes the generated audio data). The last layers of the two networks are fed into a contrastive loss function which calculates the similarity between the two audio recordings to predict the grades.

We experimented with different base networks within the Siamese architecture including a dense network, an LSTM and 1D Convolution Neural Network(CNN). There has been prior research showing the benefits of using CNNs on sequential data [27,28]. We wanted to explore their performance on our data. We report the results for 1D CNN in this work.

5.2 Multi-task Learning and Ensembling

Following the development of the decision tree, LSTM, and Siamese network models, we selected the two highest performing models across fluency and accuracy and ensembled their predictions using a simple regression model. As will be discussed in the next section, two Siamese network models (one observing textual features and the other the tonal feature set) exhibited the highest performance and were used in this process.

Our final comparison explores the usage of multitask learning [29] for the Siamese network. In this type of model, the weights of the network are optimized to predict both fluency and accuracy of speech simultaneously within a single model. Such a model may be able to take advantage of correlations between the labels to better learn distinctions between the assessment scores.

6 Results

In comparing the model results we use two measures to evaluate each model's ability to predict the Fluency and Accuracy grades: mean squared error (MSE) and Spearman correlation (Rho). A lower MSE value is indicative of superior model performance while higher Rho values are indicative of superior performance; this later metric is used to compare monotonic, though potentially non-linear relationships between the prediction and the labels (while continuous, the labels do not necessarily follow a normal distribution as students were more likely to receive higher grades).

Table 1 illustrates the model performance when comparing models utilizing each of the three described sets of features (See Sect. 4). From the table, it can be seen that in terms of MSE and Rho, the Siamese network exhibits the best performance across both metrics. It is particularly interesting to note that

Model	Features	Fluency		Accuracy	
		Rho	MSE	Rho	MSE
Siamese network	Text features	0.073	0.665	0.317	0.833
	Tonal features	0.497	0.497	-0.006	0.957
LSTM	PyAduio features	0.072	1.139	-0.066	2.868
	Tonal features	-0.096	0.648	0.042	0.929
	Text features	-0.123	1.885	0.128	3.733
Decision tree	PyAudio Features	-0.005	0.749	0.011	1.189
	Tonal features	0.285	0.649	0.107	0.960
	Text features	0.090	0.794	0.261	0.998
	Cosine similarity	0.037	0.674	0.162	0.984
	Baseline	-	0.636	-	0.932

 Table 1. Results for the different models.

Table 2. Multitasking and ensembled Siamese models

Model	Fluency		Accuracy	
	Rho	MSE	Rho	MSE
Multitasking	0.129	0.603	0.313	0.915
Ensemble (Regression)	0.477	0.490	0.34	0.839

the textual features are better at predicting the accuracy score (MSE = 0.833, Rho = 0.317), while the tonal features are better at predicting the fluency score (MSE = 0.497, Rho = 0.497).

Table 2 shows the results for the multitasking and ensemble models. We see that the Siamese model with multitasking to predict both the fluency and accuracy scores do not perform better than the individual models predicting each. This suggests that the model is not able to learn as effectively when presented with both labels in our current dataset; it is possible that such a model would either need more data or a different architecture to improve. The slight improvement in regard to Fluency MSE and Accuracy Rho exhibited by the ensemble model suggests that the learned features (i.e. the individual model predictions) are able to generalize to predict the other measure. The increase in Rho for accuracy is particularly interesting as the improvement suggests that the tonal features are similarly helpful in predicting accuracy when combined with the textual-based model.

7 Discussion and Future Work

In [5], it was found that the use of audio features helped predict fluency and accuracy scores better than a simple baseline. In this paper the textual features and tonal features explored provide even better predictive power.

A potential limitation of the current work is the scale of the data observed, and can be addressed by future research. The use of the pre-trained models may have provided additional predictive power for the tonal and textual features, but there may be additional ways to augment the audio-based features in a similar manner (i.e. either by using pre-trained models or other audio data sources). Similarly, audio augmentation methods may be utilized to help increase the size and diversity of dataset (e.g. even by simply adding random noise to samples).

Another potential limitation of the current work is in regard to the exploration of fairness among the models. It was described that the ToneNet model used training samples from men, but not women; as with any assessment tool, it is important to fully explore any potential sources of bias that exist in the input data that may be perpetuated through the model's predictions. In regard to the set of pre-trained speech recognition models provided through Google's APIs, additional performance biases may exist for speakers with different accents. Understanding the potential linguistic differences between language learners would be important in providing a feedback tool that is beneficial to a wider range of individuals. A deeper study into the fairness of our assessment models would be needed before deploying within a classroom.

As Mandarin Chinese is a tonal language, the seeming importance and benefit of including tonal features makes intuitive sense. In both the tonal and textual feature sets, a pre-trained model was utilized which may also account for the increased predictive power over the audio features alone. As all libraries and methods used in this work are openly available, the methods and results described here present opportunities to develop such techniques into assessment and feedback tools to benefit teachers and students in real classrooms; in this regard, they also hold promise in expanding to other languages or other audiobased assignments and is a planned direction of future work.

Acknowledgements. We thank multiple NSF grants (e.g., 1931523, 1940236, 1917713, 1903304, 1822830, 1759229, 1724889, 1636782, 1535428, 1440753, 1316736, 1252297, 1109483, & DRL-1031398), the US Department of Education Institute for Education Sciences (e.g., IES R305A170137, R305A170243, R305A180401, R305A120125, R305A180401, & R305C100024) and the Graduate Assistance in Areas of National Need program (e.g., P200A180088 & P200A150306), EIR, the Office of Naval Research (N00014-18-1-2768 and other from ONR), and Schmidt Futures.

References

- 1. Singer, H., Ruddell, R.B.: Theoretical models and processes of reading (1970)
- Wagner, R.K., Schatschneider, C., Phythian-Sence, C.: Beyond Decoding: The Behavioral and Biological Foundations of Reading Comprehension. Guilford Press, New York (2009)
- AlKilabi, A.S.: The place of reading comprehension in second language acquisition. J. Coll. Lang. (JCL) 31, 1–23 (2015)
- 4. Witt, S.M.: Automatic error detection in pronunciation training: where we are and where we need to go. In: Proceedings of the IS ADEPT, p. 6 (2012)

- Varatharaj, A., Botelho, A.F., Lu, X., Heffernan, N.T.: Hao Fayin: developing automated audio assessment tools for a Chinese language course. In: The Twelfth International Conference on Educational Data Mining, pp. 663–666, July 2019
- Kyriakopoulos, K., Knill, K., Gales, M.: A deep learning approach to assessing non-native pronunciation of English using phone distances. In: Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH, vol. 2018, pp. 1626–1630 (2018)
- 7. Pucher, M., Türk, A., Ajmera, J., Fecher, N.: Phonetic distance measures for speech recognition vocabulary and grammar optimization. In: Proceedings of the 41rd Congress of the Alps Adria Acoustics Association (2007)
- Bernstein, J., Cohen, M., Murveit, H., Rtischev, D., Weintraub, M.: Automatic evaluation and training in english pronunciation. In: First International Conference on Spoken Language Processing (1990)
- Litman, D., Strik, H., Lim, G.S.: Speech technologies and the assessment of second language speaking: approaches, challenges, and opportunities. Lang. Assess. Q. 15(3), 294–309 (2018)
- 10. Bernstein, J.C: Computer scoring of spoken responses. In: The Encyclopedia of Applied Linguistics (2012)
- Franco, H., Neumeyer, L., Digalakis, V., Ronen, O.: Combination of machine scores for automatic grading of pronunciation quality. Speech Commun. **30**(2–3), 121–130 (2000)
- Wei, S., Guoping, H., Yu, H., Wang, R.-H.: A new method for mispronunciation detection using support vector machine based on pronunciation space models. Speech Commun. 51(10), 896–905 (2009)
- Zechner, K., Higgins, D., Xi, X., Williamson, D.M.: Automatic scoring of nonnative spontaneous speech in tests of spoken english. Speech Commun. 51(10), 883–895 (2009)
- Ryant, N., Slaney, M., Liberman, M., Shriberg, E., Yuan, J.: Highly accurate mandarin tone classification in the absence of pitch information. Proc. Speech Prosody 7 (2014)
- Chen, C.J., Gopinath, R.A., Monkowski, M.D., Picheny, M.A., Shen, K.: New methods in continuous mandarin speech recognition. In: Fifth European Conference on Speech Communication and Technology (1997)
- 16. Loukina, A., Madnani, N., Cahill, A.: Speech- and text-driven features for automated scoring of English speaking tasks. In: Proceedings of the Workshop on Speech-Centric Natural Language Processing, pp. 67–77, Copenhagen, Denmark, September 2017. Association for Computational Linguistics
- 17. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Advances in neural information processing systems, pp. 3111–3119 (2013)
- Pennington, J., Socher, R., Manning, C.D.: Glove: global vectors for word representation. In: Empirical Methods in Natural Language Processing (EMNLP), pp. 1532–1543 (2014)
- Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. arXiv preprint arXiv:1607.04606 (2016)
- 20. Download ffmpeg. http://ffmpeg.org/download.html. Accessed 10 Feb 2019
- Giannakopoulos, T.: Pyaudioanalysis: an open-source python library for audio signal analysis. PloS one 10(12) (2015)
- 22. Zhang, A.: Speechrecognition pypi. https://pypi.org/project/SpeechRecognition/. Accessed 10 Feb 2019

- 23. Kim, J.Y., et al.: A comparison of online automatic speech recognition systems and the nonverbal responses to unintelligible speech. arXiv preprint arXiv:1904.12403 (2019)
- Koo, T.K., Li, M.Y.: A guideline of selecting and reporting intraclass correlation coefficients for reliability research. J. Chiropractic Med. 15(2), 155–163 (2016)
- Gao, Q., Sun, S., Yang, Y.: Tonenet: a CNN model of tone classification of Mandarin Chinese. Proc. Interspeech 2019, 3367–3371 (2019)
- 26. Breiman, L.: Classification and Regression Trees. Routledge, Abingdon (2017)
- Bai, S., Kolter, J.Z., Koltun, V.: An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv preprint arXiv:1803.01271 (2018)
- Huang, W., Wang, J.: Character-level convolutional network for text classification applied to chinese corpus. arXiv preprint arXiv:1611.04358 (2016)
- 29. Caruana, R.: Multitask learning. Mach. Learn. 28(1), 41–75 (1997)