
MARSYAS: a framework for audio analysis

GEORGE TZANETAKIS[†] and PERRY COOK[‡]

^{†‡}Department of Computer Science, 35 Olden Street, and [‡]Department of Music, Princeton University, Princeton, NJ 08544, USA
E-mail: gtzan@cs.princeton.edu prc@cs.princeton.edu Fax: + 1609-258-1771

Existing audio tools handle the increasing amount of computer audio data inadequately. The typical tape-recorder paradigm for audio interfaces is inflexible and time consuming, especially for large data sets. On the other hand, completely automatic audio analysis and annotation is impossible using current techniques. Alternative solutions are semi-automatic user interfaces that let users interact with sound in flexible ways based on content. This approach offers significant advantages over manual browsing, annotation and retrieval. Furthermore, it can be implemented using existing techniques for audio content analysis in restricted domains. This paper describes MARSYAS, a framework for experimenting, evaluating and integrating such techniques. As a test for the architecture, some recently proposed techniques have been implemented and tested. In addition, a new method for temporal segmentation based on audio texture is described. This method is combined with audio analysis techniques and used for hierarchical browsing, classification and annotation of audio files.

1. INTRODUCTION

There is a growing amount of audio data available on the Internet and elsewhere today. The traditional tape-recorder sample-playback paradigm for browsing, locating, manipulating and skimming audio is cumbersome and inflexible. The main reason is that it treats audio data as a linear block of samples. Traditional information retrieval (IR) (van Rijsbergen 1979), used by many of the popular Web search engines, uses computer-readable text as data and offers the ability to quickly locate and browse large amounts of data using the familiar search and 'ranked by similarity' interface. Unfortunately, there are no equivalent methods available for audio.

An obvious solution to the problem of handling large amounts of audio data is to annotate it with textual information and then use traditional IR techniques for searching. This approach works well and has the advantage of using well-known and supported techniques. On the other hand, using current interfaces, human annotation of audio is extremely time consuming.

Recently, a number of techniques for automatic analysis of audio information have been proposed (Foote 1999). These approaches work reasonably well for

restricted classes of audio. Based on these techniques, a completely automatic annotation system for audio could be envisioned. Although not impossible in theory, there are two problems with such an approach. The first is that current systems are not perfect and, therefore, annotation errors are inevitable. This problem has to do with the current state of the art, so it is possible that in the future it will be solved. There is a second problem, however, that is more subtle and not so easy to address. Audio, and especially music, is heard and described differently by each listener. There are, however, attributes of audio that most listeners will agree upon, like the general structure of the piece, the style, etc. Ideally, a system for annotation should automatically extract as much information as it can and then let the user edit and expand it.

This leads to a semi-automatic approach that combines both manual and fully automatic annotation into a flexible, practical user interface for audio manipulation. This paper describes a framework for building audio analysis tools and integrating them using a semi-automatic graphical interface. The framework has been designed to be flexible and to accommodate new algorithms easily.

In addition, a new approach to segmentation of audio files based on texture is described. The combination of temporal segmentation and sound classification significantly reduces the overhead of manual annotation and forms a powerful foundation for audio analysis applications. Moreover, it can be used to improve classification performance by using a texture-adaptive window size for integrating classification results.

1.1. Applications

There are several possible application areas for semi-automatic audio analysis tools. Digital video libraries are an active area of research that could benefit from the development of such tools. The Informedia project at Carnegie Mellon (Hauptmann and Witbrock 1997) contains a terabyte of data. Indexing the archive is done using a combination of speech-recognition, image analysis and keyword searching techniques. Audio analysis and browsing tools would enhance the current indexing techniques, especially for the regions that do not contain speech.

Detecting speech segments is very important for automatic speech recognition systems, especially when dealing with real-world multimedia data. Moreover, detecting if a speaker is male or female or determining his/her identity improves recognition performance.

More generally, audio analysis tools can be used to implement signal responsive algorithms. A compression algorithm, for example, knowing that a signal contains speech can use a method that takes advantage of speech characteristics to achieve a higher compression rate. Responding to parameters of analysed audio can also be used by interactive algorithms performing animation or sound synthesis in virtual reality simulations and computer games.

There are many libraries of sound effects and instrument samples available. Due to their large size, searching for a particular sound can be a daunting task. Using audio-similarity retrieval techniques and fast browsing can greatly accelerate this process.

1.2. Related work

A number of techniques for audio analysis have recently been proposed. In this section, some of these systems, relevant to our work, will be briefly described. A more complete overview can be found in Foote (1999).

A robust multi-feature music speech discriminator is described in Scheirer and Slaney (1997). A similar discriminator is used in Rossignol *et al.* (1998) to initially separate speech from music and then detect phonemes or notes accordingly. A multi-feature classifier based on spectral moments for recognition of steady-state instrument tones is described in Fujinaga (1998).

A retrieval-by-similarity system for isolated sounds has been developed at Muscle Fish LLC (Wold, Blum, Keislar and Wheaton 1996). Users can search for and retrieve sounds by perceptual and acoustical features, can specify classes based on these features and can ask the engine to retrieve similar or dissimilar sounds.

Speech Skimmer (Arons 1997) is an example of pushing audio interaction beyond the tape-recorder metaphor. The user can audition spoken documents at several times real-time, using time compression techniques and segmentation based on pitch. Hidden Markov Models are used in Boreczky and Wilcox (1998) and Kimber and Wilcox (1996) for segmentation and analysis of recorded meetings by speaker.

2. FRAMEWORK

All these projects use similar features, classifications and algorithms for different tasks. Therefore, in the design of our system, we made an effort to abstract the common elements and use them as architectural building blocks. This facilitates the integration of different techniques under a common framework and interface. In addition, it helps rapid prototyping since the common elements

are written once, and developing and evaluating a new technique or application requires writing only the new task-specific code.

Typically, sound analysis systems follow a bottom-up processing architecture where sensory information flows from low-level signals to higher-level cognitive representations. However, there is increased evidence that the human auditory system uses top-down as well as bottom-up information flow (Slaney 1997). A top-down (prediction-driven) approach has been used for computational auditory scene analysis (Ellis 1996). An extension to this approach with a hierarchical taxonomy of sound sources is proposed in Martin (1998). In the design of our framework, we tried to have a flexible architecture that can support these models of top-down flow and hierarchical classification as well as traditional bottom-up processing.

2.1. Architecture

The framework is named MARSYAS after a Greek mythological figure. The initials stand for MusicAl Research SYstem for Analysis and Synthesis. It is implemented using a client-server architecture. The server written in C++, contains all the signal processing and pattern recognition modules optimised for performance. The client, written in Java, contains only the user interface code and communicates with the computation engine via sockets.

This breakdown has the advantage of decoupling the interface from the computation code and allows different interfaces to be built that use the same underlying computational functionality. For example, the server can be accessed by different graphical user interfaces, scripting tools, web crawlers, etc. Both the server and the client run on Solaris, SGI, Linux and Windows (95, 98 and NT) platforms.

Special attention was given to abstracting the audio analysis process using object-oriented programming techniques. Abstract classes are used to provide a common API for the building blocks of the system, and inheritance is used to factor out common operations. The main classes of the system can roughly be divided into two categories: process-like objects and data-structure-like objects.

Process objects:

Transformations are the low-level signal processing units used by the system. They take as input a frame of sound samples and output a transformation of that frame. Some examples are: power spectral density, cepstrum, windowing, digital filtering.

Features process a frame of sound samples and output a vector which, unlike transformations, is reduced significantly in dimensionality. They typically use sound transformations for their calculations. Because

the output is a vector, several ‘physical’ features can be combined. For example, both spectral centroid and roll-off involve the calculation of power spectral density, and it is possible to bundle them together into one feature for increased performance or to use them separately for rapid prototyping.

Memories are circular buffers that hold previously calculated features for a limited time. They are used to compute means and variances of features over large windows without recomputing the features. They can have different sizes depending on the application.

Iterators break up a sound stream into frames. For each frame they use features and memories to calculate a feature vector. The result is a time-series of feature vectors which is called the feature map. Typically, there is a different iterator for each classification scheme. For example, a silence/non-silence iterator uses energy as a feature and has no memories. A more complicated iterator like the music/speech iterator uses nine features and two memories (of different size) for feature calculation.

Classifiers take as input a feature vector and return its estimated class. They are trained using labelled feature maps.

Segmentors take as input feature maps and output a signal with peaks corresponding to segmentation boundaries.

Data structure objects:

Vectors are the basic data components of the system. They are float arrays tagged with sizes. Operator overloading is used for vector operations to avoid writing many nested loops for signal processing code. The operator functions are inlined and optimised. The resulting code is easy to read and understand without compromising performance.

Sound data objects contain samples of audio as vectors with additional header information such as sample rate, channels, etc.

Feature maps are time-series of feature vectors. Feature maps can be class labelled for evaluation and training.

Time regions are time intervals tagged with annotations information.

Time lines are lists of time regions.

Time trees are arbitrary trees of time regions. They represent a hierarchical decomposition of audio into successively smaller segments (see figure 1).

All the objects contain methods to read/write them to files and transport them using the socket interface. For example, a calculated feature map can be stored and used to evaluate different classifiers without having to redo the feature calculation for each classifier.

Although the objects form a natural bottom-up hierarchy, top-down flow of information can be expressed

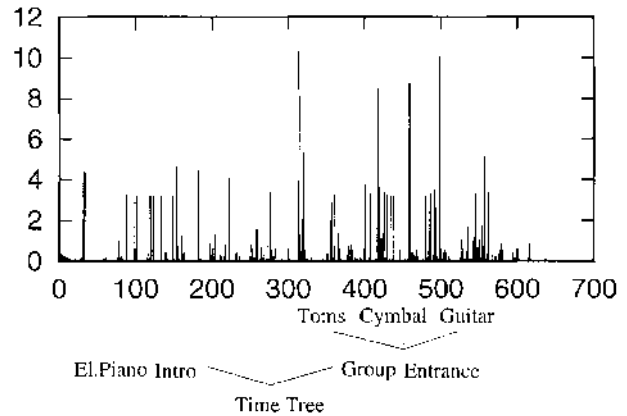


Figure 1. Segmentation peaks and corresponding time tree.

in the framework. As a simple example, a silence feature can be used by an iterator for music/speech to avoid calculating features on silent frames. Similarly, hierarchical classification can be expressed using multiple iterators with different features.

2.2. Features

In our system we have implemented a number of the features that have been proposed in the literature. These features form a pool from which different algorithms can pick specific features depending on the specific task. Many classification algorithms, for example, must be invariant to loudness so they would not use energy as a feature. On the other hand, a segmentation algorithm would probably include energy as one of the features indicating a change in texture.

Some of the currently supported features, with references to systems that describe and use them, are given below.

Spectral centroid is the balancing point of the spectrum. It can be calculated using

$$(1) \quad C = \frac{\sum_i i A_i}{\sum_i A_i},$$

where A_i is the amplitude of frequency bin i of the spectrum (Wold *et al.* 1996, Scheirer and Slaney 1997, Fujinaga 1998, Rossignol *et al.* 1998).

Spectral moments are statistical measures that characterise the shape of the spectrum (Fujinaga 1998).

Spectral flux is the 2-norm of the difference between the magnitude of the short time Fourier transform (STFT) spectrum evaluated at two successive sound frames. The STFT is normalised in energy (Scheirer and Slaney 1997, Rossignol *et al.* 1998).

Pitch is a pitch estimate for the frame and can be calculated using various different techniques (Rabiner, Cheng, Rosenberg and McGonegal 1976).

Harmonicity is a measure of how strong the pitch perception for a sound is (Wold *et al.* 1996). It can also be used for voiced/unvoiced detection.

Mel-frequency cepstral coefficients (MFCC) are commonly used in speech recognition (Kimber and Wilcox 1996, Scheirer and Slaney 1997). They are a perceptually motivated compact representation of the spectrum (Hunt, Lennig and Mermelstein 1996).

Linear prediction (LPC) reflection coefficients are used in speech research as an estimate of the speech vocal tract filter (Makhoul 1975).

Other features supported include zero crossings, RMS, spectral roll-off and others. For all these features, means, variances and higher-order statistics over larger time windows can be calculated using memories. New features can easily be added to the system by writing only the code for computing the feature value from a frame of sound samples.

2.3. Classifiers

Currently, two statistical pattern recognition classifiers are implemented as part of the system. For a more complete description of these classifiers and statistical pattern recognition in general, refer to Duda and Hart (1973).

The Gaussian (MAP) classifier assumes each class can be represented as a multi-dimensional normal distribution in feature space. A labelled data set is used to train the classifier by calculating the parameters for each particular class. This classifier is typical of parametric statistical classifiers that assume a particular form for the class probability density functions.

Unlike parametric classifiers, the K-nearest-neighbour classifier directly uses the training set for classification without assuming any mathematical form for the underlying class probability density functions. Each sample is classified according to the class of its nearest neighbour in the training data set. In K-NN, the K nearest-neighbours are calculated and voting is used to determine the class.

Due to the flexibility of the architecture, new and more advanced classifiers can easily be added. Gaussian mixture and neural network classifiers are currently under development. The ability to have different classifiers allows trade-off between classification speed and accuracy, depending on the application.

2.4. Framework evaluation

The data used for evaluating the system consists of about two hours of audio data. There are forty-five minutes of speech, forty-five minutes of music, and about thirty minutes of mixed audio. Radio, live recordings of speech, compact discs and movies representing a variety of speakers and music styles were used as data sources.

As a test for the architecture, a music/speech discriminator similar to the one described in Scheirer and Slaney (1997) and Rossignol *et al.* (1998) was implemented. The implementation was used to test and refine the design of the framework. The discriminator runs in real time on an SGI 02 workstation. Figure 2 shows the layout and information flow of the implementation.

The performance of our system for music/speech discrimination is comparable to the recognition accuracy of current systems (Scheirer and Slaney 1997, Rossignol *et al.* 1998) (90.1% for the K-NN(5) classifier). A direct comparison is impossible due to the differences in data sets and classifiers. We use a cross-validation testing framework (Scheirer and Slaney 1997) to ensure that the evaluation is not dependent on the particular test and training sets we have used.

In addition, an instrument identification system for steady-state tones similar to Fujinaga (1998) has been implemented. Similarity retrieval as described in Wold *et al.* (1996) is also supported.

3. SEGMENTATION

Segmentation is based on the idea that transitions in audio texture will result in sudden changes of values in most of the features. This idea is consistent with research in psychoacoustics (Bregman 1990) where multiple perceptual cues are combined to indicate the onset of a new sound event.

The algorithm works in four stages:

- (1) A time series of feature vectors V_i is calculated by iterating over the sound file. Each feature vector can be thought of as a short description of the corresponding time-frame.
- (2) A distance-metric $\Delta_i = \|V_i - V_{i-1}\|$ is calculated between successive frames of sound. In our implementation we use a Mahalanobis distance. It is defined by

$$(2) \quad D(x,y) = (x - y)^T \Sigma^{-1} (x - y),$$

where Σ is the feature covariance matrix calculated from the training set. Other distance metrics, possibly using relative feature weighting, can also be used.

- (3) The derivative $d\Delta_i/dt$ of the distance signal is taken. Thresholding is then used for finding the peaks of the result. The derivative of the distance will be low for slowly changing textures and high during sudden transitions. Therefore, the peaks roughly correspond to texture changes.
- (4) Peaks are picked using simple heuristics and used to create the segmentation of the signal into time regions. As a heuristic example, a minimum duration between successive peaks can be set to avoid small regions. The result is stored as a time line (i.e.

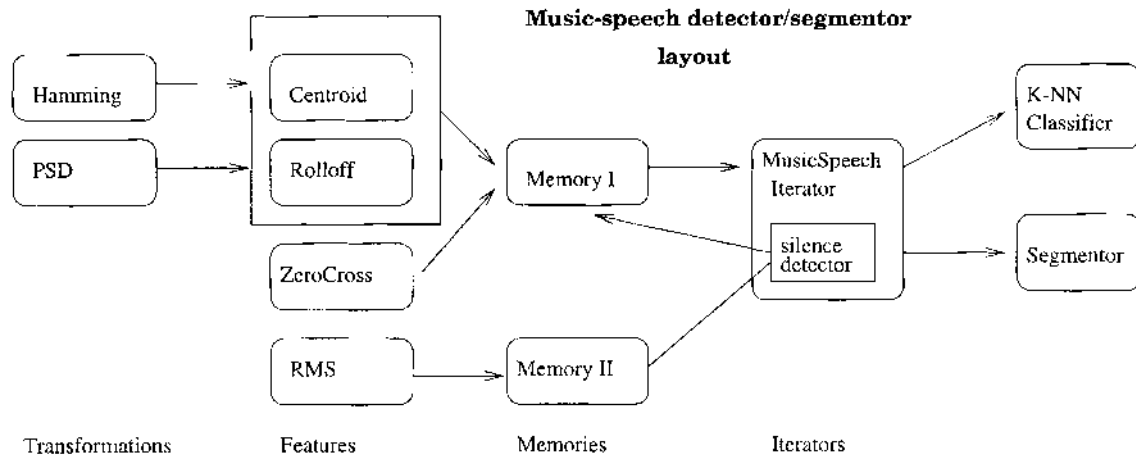


Figure 2. Music speech detector.

a list of time intervals) and can be used for browsing and annotating (see figure 1).

3.1. Combining segmentation and classification

Most of the classification methods proposed in the literature report improved performance if the classification results are integrated over larger time windows. However, using fixed size integration windows blurs the transition edges between classes. Usually, test data consist of files that do not contain transitions to simplify the evaluation; therefore this problem does not show up. In real-world data, however, transitions exist and it is important to preserve them.

The described segmentation method provides a natural way of breaking up the data into regions based on texture. These regions can then be used to integrate classification results. That way, sharp transitions are preserved and the classification performance is improved because of the integration. Initial experiments in a number of different sound files confirm this fact. A more detailed quantitative evaluation of how this method compares to fixed-window integration is under way.

3.2. Segmentation results

The method has been tested using various sound files from our data set. Representative examples include detecting a guitar solo entrance or a cymbal crash, change of speaker, transitions from music to speech, and musical structure for cyclic (ABAB type) pop songs.

The segmentation results are difficult to evaluate quantitatively because they depend on the choice of features and parameter values like the memory-size, peak-threshold, distance-metric, and others. Typically, the resulting regions are perceptually meaningful. For now, fine-tuning of the parameters is done by the user (or the programmer) depending on the desired result. For example, lowering the peak-threshold results in more

and smaller regions. Automatic parameter adjustment is investigated.

4. USER INTERFACE

The interface looks like a typical tape-recorder style waveform-editor. However, in addition to the typical play, fast-forward, rewind and stop buttons, it allows skipping by either user-defined fixed duration blocks or time lines containing regions of variable duration. These time lines are created either by hand or automatically using the segmentation method described above.

Skipping and annotating using regions is much faster than manual annotation, in the same way that finding a song on a CD is much faster than finding it on a tape.

The user can select a region and retrieve similar sounds. Another possibility is to classify the region using one of the available classification schemas like the music/speech discriminator. Finally, each time region can be annotated by multiple keywords.

In addition, the user can combine time regions to form a time tree that can be used for multi-resolution browsing and annotation. The tree captures the hierarchical nature of music pieces, and can therefore be used for musical analysis.

5. APPLICATIONS

The initial description of MARSYAS appeared in Tzanetakis and Cook (1999a). A more detailed description of the segmentation algorithm as well as some user experiments performed for evaluation are given in Tzanetakis and Cook (1999b). A number of different applications have been developed using our system. Most of them are undergraduate senior theses and class projects in Princeton. Some examples are an instrument family classification system and a content-based music library system. In addition, we have investigated the use of MPEG audio compressed data as a basis for feature calculation. Finally, the user interface has been used to

conduct user experiments on how humans annotate and segment audio (Tzanetakis and Cook 2000).

6. FUTURE WORK

On the interface side, we plan to support multiple time lines and fast keyword search in annotations. Another interesting application of our segmentation scheme is audio thumbnails. For each region, a characteristic segment has to be selected. These segments are then used to create a shorter summary version of the original soundfile.

On the computational side, we are investigating using a more perceptually accurate front end to the system. A number of computational models of the ear physiology have been proposed in the literature (Slaney and Lyon 1990, 1993) and can be used as a basis for feature calculation.

Most attempts to build music analysis systems in the past have tried to extract content by first transcribing the music into symbolic notation and then using music theory to characterise it. This approach has been challenged by Scheirer (1996) and Martin, Scheirer and Vercoe (1998). Current systems try to analyse structure and content directly from features calculated from the audio signal. Such systems can easily be implemented and evaluated using our framework. As a simple example of music analysis, the structure of cyclic pop songs can be revealed using our segmentation scheme. We believe that a combination of our segmentation scheme with beat tracking methods (Scheirer 1998) can offer significant information for music style identification and music analysis.

Finally, we plan to write a Web crawler that will automatically create segmentation time lines and annotations for files on the Web.

7. SUMMARY

We designed and implemented a flexible framework for building and integrating audio analysis tools. A number of existing techniques were implemented and tested to evaluate the framework. A new method for audio segmentation based on texture was presented. This method combined with the analysis tools and using a semi-automatic user interface offers significant improvements for audio searching, annotation and browsing. In addition, it can be used to improve classification performance by using the results of the segmentation to adaptively change the classification integration window size.

REFERENCES

Arons, B. 1997. Speechskimmer: A system for interactively skimming recorded speech. *ACM Trans. Computer Human Interaction* **4**: 3–38.

- Boreczky, J., and Wilcox, L. 1998. A hidden Markov model framework for video segmentation using audio and image features. *Proc. Int. Conf. on Acoustics, Speech and Signal Processing*, Vol. 6, pp. 3,741–4.
- Bregman, A. 1990. *Auditory Scene Analysis*. Cambridge, MA: MIT Press.
- Duda, R., and Hart, P. 1973. *Pattern Classification and Scene Analysis*. New York: John Wiley & Sons.
- Ellis, D. 1996. *Prediction-driven Computational Auditory Scene Analysis*. PhD Thesis, MIT Dept of Electrical Engineering and Computer Science.
- Foote, J. 1999. An overview of audio information retrieval. *ACM Multimedia Systems* **7**: 2–10.
- Fujinaga, I. 1998. Machine recognition of timbre using steady-state tone of acoustic instruments. *Proc. ICMC 98*, pp. 207–10.
- Hauptmann, A., and Witbrock, M. 1997. Informedia: News-on-demand multimedia information acquisition and retrieval. *Intelligent Multimedia Information Retrieval*, Chap. 10, pp. 215–40. Cambridge, MA: MIT Press. <http://www.cs.cmu.edu/afs/cs/user/alex/www/>
- Hunt, M., Lennig, M., and Mermelstein, P. 1996. Experiments in syllable-based recognition of continuous speech. *Proc. 1996 ICASSP*, pp. 880–3.
- Kimber, D., and Wilcox, L. 1996. Acoustic segmentation for audio browsers. *Proc. Interface Conference*. Sydney, Australia.
- Makhoul, J. 1975. Linear prediction: A tutorial overview. *Proc. IEEE* **63**: 561–80.
- Martin, K. 1998. Toward automatic sound source recognition: identifying musical instruments. *NATO Computational Hearing Advanced Study Institute*. Il Ciocco IT.
- Martin, K., Scheirer, E., and Vercoe, B. 1998. Musical content analysis through models of audition. *Proc. ACM Multimedia Workshop on Content-Based Processing of Music*. Bristol, UK.
- Rabiner, L., Cheng, M., Rosenberg, A., and McGonegal, C. 1976. A comparative performance study of several pitch detection algorithms. *IEEE Trans. on Acoustics, Speech and Signal Processing ASSP-24*: 399–417.
- Rossignol, S., Rodel, X., Soumagne, J., Colette, J.-L., and Depalle, P. 1998. Features extraction and temporal segmentation of acoustic signals. *Proc. ICMC 98*, pp. 199–202.
- Scheirer, E. 1996. Bregman's chimerae: Music perception as auditory scene analysis. *Proc. Int. Conf. on Music Perception and Cognition*. Montreal.
- Scheirer, E. 1998. Tempo and beat analysis of acoustic musical signals. *Journal of the Acoustical Society of America* **103**(1): 588, 601.
- Scheirer, E., and Slaney, M. 1997. Construction and evaluation of a robust multifeature speech/music discriminator. *IEEE Trans. on Acoustics, Speech and Signal Processing (ICASSP'97)*, pp. 1,331–4.
- Slaney, M. 1997. A critique of pure audition. In D. Rosenthal and H. Okuno (eds.) *Computational Auditory Scene Analysis*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Slaney, M., and Lyon, R. 1990. A perceptual pitch detector. *Proc. of the 1990 Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 357–60. Albuquerque, NM. IEEE.

- Slaney, M., and Lyon, R. 1993. On the importance of time – a temporal representation of sound. In M. Cooke, B. Beet and K. Crawford (eds.) *Visual Representations of Speech Signals*, pp. 95–116. New York: John Wiley & Sons Ltd.
- Tzanetakis, G., and Cook, P. 1999a. A framework for audio analysis based on classification and temporal segmentation. *Proc. 25th Euromicro Conf. Workshop on Music Technology and Audio Processing*. Milan, Italy. IEEE Computer Society.
- Tzanetakis, G., and Cook, P. 1999b. Multifeature audio segmentation for browsing and annotation. *Proc. 1999 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, WASPAA99*. New Paltz, NY.
- Tzanetakis, G., and Cook, P. 2000. Experiments in computer-assisted annotation of audio. *Proc. ICAD2000*. Atlanta (to be published).
- van Rijsbergen, C. 1979. *Information retrieval*, 2nd edn. London: Butterworths.
- Wold, E., Blum, T., Keislar, D., and Wheaton, J. 1996. Content-based classification, search and retrieval of audio. *IEEE Multi-media* 3(2): 27–36.