Journal of the American Society for Information Science, 2003

Popular Music Access: The Sony Music Browser

François Pachet, Amaury Laburthe, Aymeric Zils, Jean-Julien Aucouturier

Sony CSL-Paris, pachet@csl.sony.fr

Abstract:

The intentionally ambiguous expression "Popular Music Browser" reflects the two main goals of this project, which started in 1998, at Sony CSL laboratory. First, we are interested in human-centered issues related to browsing "Popular Music". Popular here means that the music accessed to is widely distributed, and known to many listeners. Second, we consider "popular browsing" of music, i.e. making music accessible to non specialists (music lovers), and allowing sharing of musical tastes and information within communities, departing from the usual, single user view of digital libraries. This research project covers all areas of the music-to-listener chain, from music description - descriptor extraction from the music signal, or data mining techniques -, similarity based access and novel music retrieval methods such as automatic sequence generation, and user interface issues. This paper describes the scientific and technical issues at stake, and the results obtained, and is illustrated by prototypes developed within the European IST project *Cuidado*.

Introduction

The Route to EMD

Electronic Music Distribution usually refers to the technical issues of transporting music data across networks, copy protection and copyrights management. These enabling technologies are necessary ingredients for making more music accessible to more people. However, there is much more to EMD than telecommunication and protection. In our view, the basic role of EMD is to make possible the shift from a mass-market approach to music distribution to a 1-to-1 distribution approach, and therefore to provide some sort of digital mapping between music and people. What, at first glance, looks like a marketing issue has, however, deep ramifications in many scientific domains, and the realization of this digital map requires new fundamental results in domains far outside marketing, for several reasons.

First, size. Estimations based on major label catalogues yield a total of 10 million music titles; restricting ourselves to only published, occidental, popular music (including Rock, Jazz, Classical and many other genres, but excluding so-called “Indies”, and traditional, culture-dependent productions). Second, the number of people connected or to be connected to modern telecommunication networks is of the order of a several hundred millions. Linking music to people has traditionally been performed through so-called mass-market schemes, that is by distributing only a small fraction of music hits to a large fraction of people. Times are changing as we know, and today, more personalized distribution schemes are both possible and desired by the vast majority of users, by the artists and the record companies.

Second, music access is an ill-defined problem: we don’t want to simply “access” music, do we? Music selection and listening is not an inherently rational process, as compared to, e.g. looking for bibliographical references on the Internet. Not only the users do not always know how to specify what they look for (the *language mismatch problem* created by domain ontologies that users do not understand, is addressed e.g. by Belkin, 2000). But users do not always even *know* what they look for. We can expect a reasonable EMD system to at least know about this situation, and not systematically force the user into a systematic query/answer protocol. The design of such an EMD system requires that we know more about what users want to do with music. Unfortunately, music listening is a badly understood phenomenon today, and studies and investigations in the field of the psychology of popular music are only just starting. We will show promising avenues of research and techniques which address this issue of modelling the user goals - or lack thereof.

Assumptions on users

The very idea of Popular Music Browsing implies several assumptions concerning our typical user. First, this user is a music lover of some sort, interested in music, and ready to spend time searching in a music catalogue. Music browsing, and particularly content-based access browsing, is probably not suited to users only marginally interested in music, as these needs are perfectly addressed by many other schemes, radio for instance. Second, we assume that the user knows only a small fraction of the music catalogue. For large catalogues, this assumption often holds, including for personal music collections, in particular collections obtained from peer-to peer systems: it is virtually impossible to "know" and remember tens of thousands songs. Conversely, content-based access does not make sense for smaller size collections that are perfectly known by their user.

Finally and most importantly, we assume that the music in the catalogue is potentially "interesting" for users, or at least, that the catalogue contains many titles that the user does not know, but would eventually like, if he was exposed to them. The ultimate and ambitious goal of the Popular Music Browser (MB) is indeed to produce so-called "Aha" experiences in users, experiencing the joy of discovering something that fits his taste, however badly defined or expressed, rather than fetching exactly titles corresponding to precise queries. This last point is particularly important because there are many situations in which users do not wish to go through such experiences, and prefer - or have to use – more rational, well-defined search procedures.

Existing Popular Music Access Systems

There are today (in 2003) many on line searchable music databases. Among the many systems available, we can quote the following most interesting proposals.

First, purely editorial systems propose systematic editorial information on popular music, including albums track listings (CDDB), information on artists and songs (AMG and Muze). This information is created by music experts, or in a collaborative fashion (CDDB). These systems provide useful services for EMD systems, but cannot be considered as fully-fledged EMD systems *per se*, as they provide only superficial and incomplete information on music titles, supposed to exist somewhere else.

The MoodLogic browser propose a complete solution for Popular Music access. The core idea of MoodLogic is to associate metadata to songs automatically thanks to two basic techniques: 1) an audio fingerprinting technology able to recognize music titles on personal hard disks, and 2) a database collecting user ratings on songs, which is incremented automatically, and in a collaborative fashion. An ingenious proactive strategy is enforced to encourage users to rate songs, in order to get tokens that allow them to get more metadata from the server. Moodlogic relies entirely on metadata obtained from user ratings and does not perform any acoustic analysis of songs.

However, collaborative music rating does not exhaust the description potential of music, and our Browser proposes many other types of metadata.

Other proposals have been made either for fully-fledged music browsers, or for ingredients to be used in browsers (fingerprinting techniques, collaborative filtering systems, metadata repositories, e.g. Wold et al. 1996) that we cannot cover here for reasons of space. We will describe in this paper only the parts of our project that we think are original and may contribute to address the needs of our targeted users.

Editorial Information

The MB aims at exploiting all possible metadata that can be extracted or accumulated for music titles. We only limit ourselves to metadata that is not already available through standardized procedures. This includes information on dates of creation, liner notes, etc. that can be obtained through AMG, Muze and the likes.

Artists as identifiers and holders of editorial metadata

Artists (taken in the most general sense) are key *music identifiers* for many users: Yesterday is by “The Beatles”, and “The 5th symphony” is by Beethoven. Artists are used also for solving ambiguity: “With a Little Help from my Friends” by the Beatles, is definitely not the same tune as the version by Bruce Springsteen. The “Stabat Matter” by Pergolese is not the one by Boccherini, etc. We call these artists “primary artists” as they are most commonly used to identify music titles. These examples show that primary artists are common ways of identifying music titles but also that the role of primary artists changes with styles: in Classical music, primary artists are usually composers. In non Classical music they are usually performers. In our Browser, we introduced the notion of primary artists in a deliberate ambiguous way, to cope for Classical and non Classical music in a uniform way.

There are cases where primary artists are not enough for characterizing the identity of a piece. The “1st partita” of Bach has been recorded by Glenn Gould, and also by many other pianists, and this distinction is of course very important: not only for interpreters, but also for conductors (for orchestral pieces). In non-Classical music the need for secondary artists is also obvious, for instance to indicate that the Springsteen version of “A little help” is indeed a Beatles song.

Existing repositories of editorial information do not provide systematic schemes for accessing artists and their relations to songs. This led us to constitute a database of artists, or more generally of “Musical Human Entities” (MHE), including both performers, composers, but also groups (the Beatles), orchestra (the Berlin Philharmonic), duets (Paul McCartney & Michael Jacskon). To each artist (or MHE) is associated a limited but useful set of properties in fixed ontologies: type (composer, singer, instrumentist, etc.), country of origin, language (for singers), type of voice (for singers also), main instrument (for instrumentists). Other information concern the relation MHE entertain with each other. For instance, Paul McCartney is a *memberOf* The Beatles, and artist Phil Collins a *memberOf* the group Genesis. This information is easy to enter and encode, and provides very useful linking possibility for query by similarity (see below). The Editorial MHE database may be seen more as a knowledge base than a database. Figure 1 shows a snapshot of the Editorial Information Manager, where this information is shown.

<FIGURE 1 Editorial Data Manager> ABOUT HERE

Song Identification

The issue of audio music identification has receive tremendous attention in the last years, due to the illegal proliferation of mp3 files on the Internet. Many methods have been proposed to identify music signals automatically, in particular audio based on fingerprinting (Allamanche, 2001). These techniques are sufficiently robust to work well on large databases, but rely on the existence of huge databases of audio fingerprints. Other, less sophisticated techniques can be used to identify music titles coming form the Internet. We proposed in (Pachet & Laigre, 2001) a technique for parsing file names of arbitrary syntax. Mp3 files are often given ID3 tags, which often contain enough information to identify primary artists and song title information.

In the MB, we follow a compromise approach for title identification. A central server (see Section on architecture) contains editorial information about artist and songs, as well as rudimentary acoustic signal information (duration, normalized energy). The combined used of textual and basic acoustic information is sufficient in most cases to identify songs properly. An approximate string matching mechanism is used to propose users with the most probable matches, and this techniques proves extremely efficient in practice, with the advantage of not requiring complex audio signal processing nor huge fingerprinting database.

Genre

The issue of musical genre taxonomies is an old and tricky one. Genres are badly needed for accessing music, and are as badly ill-defined. Our studies on existing taxonomies of genres have shown that there is no consensus, and that a consensus is probably impossible (Aucouturier & Pachet, 2003). However, genres are needed, and after several years of trials and errors, we ended up with a simple two-level genre taxonomy consisting of 250 genres. The main property of this taxonomy is flexibility: users can classify artists or songs either in a generic way (Classical, Jazz), more precisely (Jazz / BeBop, Classical/Baroque). However, simpler taxonomies may also produce frustration, as some categories may contain artists or songs that users would consider very different. To make our taxonomy more flexible, we have introduced an optional “keyword” field, which may contain free words. These words may be entered by users to further refine their own classification perspective on artists or songs. This simple yet flexible approach has the advantage of uniformity: artists and songs are classified in the same taxonomy, allowing for various degrees of precision. For instance, The Beatles is classified in “Pop / Brit”, but Beatles songs may be classified in other genres (e.g. “Revolution 9” is “Rock / Experimental”).

Most importantly, the database of artists and songs is made public and shareable by all users. A mechanism allows users to choose between their own private view on genres, and the view most commonly shared (see Section on Architecture).

Core metadata extraction techniques

The main type of metadata that the MB proposes for songs besides editorial information is *acoustic* metadata, i.e. information extracted from the audio signal. The Mpeg7 standard aims at providing a format for representing these information, and a specialized audio group produces specific constructs to represent musical metadata (Herrera et al., 1999).

However, music metadata in Mpeg7 refers in general to low-level, objective information that can be extracted automatically in a systematic way. Typical descriptors (called LLD for Low-Level Descriptors in the Mpeg7 jargon) proposed by Mpeg7 concern superficial signal characteristics such as Means and Variance of amplitude, spectral frequencies, spectral centroid, ZCR (Zero Crossing Rate), etc.

Concerning high-level descriptors that can be mapped to high-level perceptual categories, Mpeg7 is strictly concerned with the format for representing this information, and not the extraction process *per se*.

Extracting High-Level Music Percepts

We have conducted in the project several studies focusing on particular dimensions of music that are relevant in our context. In particular, we have proposed a rhythm extractor (Zils & al. 2002), that is able to extract the time series of percussive sounds in music signals of popular music. Rhythm information is a useful extension of tempo or beat, as proposed e.g. by (Sheirer, 1998). In particular, careful studies of popular music signals show that rhythm structure may be very complex, even for very popular tunes: the band “Earth wind & Fire” has produced popular songs (e.g. “September”) usually considered by musicologists as vulgar, but which contain in effect extremely complex rhythm pattern and structure.

However, many things remain to be done in the field of rhythm. One key issue seems to rely not so much in how to extract rhythm, but how to exploit the information: most people are unable to describe rhythm with words, and even less to produce rhythm (our attempts at designing a query by rhythm did not prove successful).

We have also addressed other dimensions of music pertaining to poplar music access, including perceptual energy (the difference in perceived intensity between Hard rock titles and acoustic folk ballads), and global timbre color. All these descriptors as well as the editorial descriptors can be used for making specific queries to the database (see Figure 2). This latter descriptor showed particularly interesting, as it allows to match titles of possible very different genres based solely on timbral color information, and we have shown that such similarity measures are key steps for achieving Aha effects (Aucouturier & Pachet, 2002a).

<FIGURE 2 Photo Query Panel > IS ABOUT HERE

EDS: A General Framework for Extracting Extractors

These various studies in descriptor extraction from acoustic signals have shown that the design of an efficient acoustic extractor is a very heuristic process, which requires sophisticated knowledge of signal processing, intuitions, and experience. Indeed, most approaches in feature extraction as published in the literature consist in using statistical analysis tools to explore spaces of combinations of LLD. The approaches proposed by (Peeters et Rodet, 2002; Sheirer and Slaney, 1997) typical fall in this category. However, these approaches are not capable of yielding precise extractors, and depend on the nature of the palette of LLD, which usually do not capture the relevant, often intricate and hidden characteristics of audio signals. Consequently, designing extractors is very expensive and hazardous.

These experiments have given rise to a systematic approach to feature extraction, embodied in the EDS system (Pachet et Zils, 2003). Departing from the usual LLD approach, the idea of EDS is to automate – in part or totally – the process of designing extractors. EDS searches in a richer and more complex space of signal processing functions, much in the same way than experts do: by inventing functions, computing them on test databases, and modifying them until good results are obtained.

To reach this goal EDS uses a genetic programming engine, augmented with fine grained typing system, which allows to characterize precisely the inputs and outputs of functions. EDS uses also rewriting rules to simplify complex signal processing functions (see the example of the Perceval equality being used by EDS to simplify the expression in Figure 3). Finally EDS uses expert knowledge to guide its search, in the form of heuristics. Typical heuristics include “do not try functions which contain too many repetition of the same operator”, or “apply twice a fft on a signal is interesting, but not 3 times", or also "spectral coefficients are particularly useful when applied on signals in the temporal domain, possible filtered", etc.

<FIGURE 3 PHOTO EDS> IS ABOUT HERE

The current extractors targeted by EDS are Perceptual Energy (or a refinement of the descriptor we designed by hand), discrimination between songs and instrumental, discrimination between studio and live versions of songs, harmonic complexity, etc. The ambitious goal of EDS makes it a project in itself, as it aims at capturing complex knowledge, in an expanding field. However, we think that the contribution to the MIR community is potentially important as it is a first step towards a unified vision of high level audio feature extraction. Here also, some sort of collaborative effort is needed, as good heuristics typically come from experienced users, and must be confronted to many situations to be useful and a coming online version of EDS should allow signal processing experts to test and share their knowledge in a collaborative way.

Similarity

The notion of similarity is of utter importance in the field of music information retrieval, and the expectation to have systems that find songs that are "similar" to one or several seed songs is now second nature. However, here again, similarity is ill-defined, and it can be of many different sorts. For instance, one may consider all the titles by a given artist as similar. And they are, of course, artist-wise. Similarity can also occur at the feature level. For instance, one may consider that Jazz saxophone titles are all similar. Music similarity can yet occur at a larger level, and concern songs in their entirety. For instance, one may consider Beatles titles as similar to titles from, say, the Beach Boys, because they were recorded in the same period, or are considered as the same “style”. Or two titles may be considered similar by a user or a community of users for no objective reason, simply because they think so.

The MB proposes two main forms of similarity: feature-based and cultural. Feature-based similarity is trivially obtained by defining similarity measures from the metadata obtained and described above, either editorial or acoustic. As we stated above, timbral similarity proved very useful, but most descriptors yield implicitly similarity measures (e.g. similarity of tempo, of energy, or similarity based on artist relationships, etc.) that can also be useful in some circumstances.

Cultural similarity refers to global similarity between music titles, that are produced by analysing patterns in communities of users. These kinds of similarity are non objective (or rather, non signal-based) but may be extracted automatically, at least to some extent, by data mining techniques. The most well known data mining technique is collaborative filtering and has been long applied to music recommendation systems (Cohen & Fan, 2000). Collaborative filtering consists in analysing relation between user profiles. Profiles are typically obtained by observing user behaviour, such as buying in CD retail sites or repeated selection in music players. Other sorts of data mining technique can be used to produce similarity, in particular through the use of co-occurrence techniques, applied to textual corpuses. In this latter case, two titles are considered similar if they appear repeatedly on the same web pages, or textual document. We have conducted a series of experiment in this spirit (Pachet et al. 2001), showing that the extracted similarities are usually very relevant.

The Music Browser project contains a sub project consisting in designing a web crawler and textual co-occurrence analyser to produce such similarities on a large scale. Yet other similarity relations can be extracted by studying more specific corpuses of textual information such as radio playlists. In this case, co-occurrence is defined by titles being played consecutively in a given program. The same kind of similarity measure can be defined, but the results are usually quite different. In particular, radio stations tend to produce similarities which are more metaphorical, and therefore possibly more interesting (and less obvious).

An interesting issue resulting from these studies is the comparison between these different sorts of similarity. For instance, a starting title such as “Le moment de vérité” played by Ahmad Jamal, is considered by the MB as similar timbre-wise to “Humoresque Op. 20” by Schumann or “Blue and sentimental” by Hank Jones, but culturally, it is closer to “Ahmad’s blues” by Miles Davis, because of the strong relationship between these two players, captured by the web crawler. Of course, there is no grounded truth here, and all these similarities are relevant. The next issue to solve is to aggregate these similarities, or at least propose users simple and meaningful ways of exploiting these different techniques. Much remains to be done here also and this is one of our privileged direction of study. Figure 4 shows the similarity panel of the MB, in which users may select freely the similarity distance.

<FIGURE 4 Similarity Panel> IS ABOUT HERE

User Interfaces ?

We have covered so far the core technologies for producing content descriptions of music titles. A key issue is the exploitation of these information on the user side. The graphical interface issue is problematic because of the great variety of behaviours of users, and because the actual devices that will be used for large scale access to music catalogues are still unknown (computers ? set-top boxes ? PDAs ? telephones ? Hard-disk Hi-FI ?). Many user interfaces have been proposed for music access systems, from straightforward feature-based search systems (MongoMusic) to innovative graphical representations of play lists. For instance, Gigabeat display music titles in spirals to reflect similarity relations titles entertain with each other. The gravitational model of SmartTuner of mzz.com, represent titles as mercury balls moving graciously on the screen, to or from “attractors” representing the descriptors selected by the user. However gracious, these interfaces impose a fixed interaction model, and assume a constant attitude of users regarding exploration: either non-explorative - music databases in which you get exactly what you query - or very exploratory. But the users may not choose between the two, even less adjust this dimension to their wish. The current interface of the Music Browser aims at allowing users to choose between many modes of music access: explorative, precise, focused or hazardous. An original feature introduced by the Browser to this aim is a powerful playlist generation system, based on constraint satisfaction techniques (Pachet et al. 2000 ; Aucouturier & Pachet, 2002b). This technique allows user to get entire music playlists from a catalogue, by specifying only abstract properties on the playlist, such as continuity (progressive increase in tempo), distribution (no more than 2 titles by the same artist in a row) or cardinality (30 % of Rock titles). However, the ultimate music access interface that would allow all these interactions mode yet remain simple is yet to be invented.

Architecture: private versus shared information

A final point of the Music Browser project deserves some attention as it relates to the sharing of music information with other, and the fostering of musical communities. The great success of peer-to-peer systems (Napster, Kazaa and their descendants) has shown that many users wished to access music beyond their own catalogues. However exciting, these systems suffer from a great lack of search features: most search is text based, and furthermore, the texts associated to songs are no controlled, and therefore the search may turn to be very hazardous and tedious.

The Music Browser is based on a client server architecture that allows user to share: not songs (this is illegal!), but rather musical metadata. This idea was introduced by MoodLogic, whose browser is already based on such an architecture: when a user adds titles in the browser, these titles are recognized (by a request to a server containing thousands of fingerprints). Then the musical metadata is fetched and loaded in the local user database.

The same kind of architecture has been designed for our browser. The difference lies in the nature of the metadata being transferred (in our case both editorial and acoustic descriptors are fetched from the central server). Furthermore, we exploit the architecture to decentralize the computation of acoustic metadata. Indeed, the computation of acoustic descriptors is still a costly process. Clients application compute only the metadata of titles which are not recognized by the server. Once computed, the metadata is sent to the server, so that the next user requesting the same song will directly get the associated metadata.

Technically the server is implemented as a MySQL database, and a PhP server for artist and song editorial information. A Servlet server handles services requiring non trivial computation such as approximate string matching queries and metadata synchronization.

Conclusion

We have described the Popular Music Browser project at CSL, by focusing on the and the technical issues related to content-based extraction and access methods. The end result is designed to account for a wide range of listening behaviour, and for users interested in discovering unknown titles, rather than for librarians who know exactly what they look for. This sort of music access is of course not the only one, but we believe that addressing this ill-defined issue of "looking for music" with content-based techniques is of utter importance in our society of insatiable quest for entertainment. Many issues have already been addressed and some with successful results. Many other issues remain to be addressed and will probably pop up when music listeners get used to these novel approaches to browsing. The fun is only beginning.

References

Allamanche, E. Herre, J. Helmuth, O. Frba, B. Kasten, T. and Cremer, M. (2001) "*Content-Based Identification of Audio Material Using MPEG-7 Low Level Description*" in Proc. of the Int. Symp. of Music Information Retrieval.

Aucouturier, J.-J. Pachet, F. (2002a) Music similarity measures: what's the use ? Proc. of the 3rd Int. Symp. of Music Information Retrieval (ISMIR02), Paris, October.

Aucouturier, J.-J. Pachet, F. (2002b) Scaling up Playlist generation, In Proc. of the IEEE Int. Conf. on Multimedia and Expo (ICME 02), Lauzanne, October.

Aucouturier, J.-J. Pachet, F. (2003) Musical Genre: a Survey, In Journal of New Music Research, 32:1.

Belkin, N. (2000) Helping people find what they don't know. *CACM* Vol. 43, N. 8, August 2000, pp. 58-61.

Cohen, W., Fan, W. (2000) Web-Collaborative Filtering: Recommending Music by Crawling The Web, 9th International World Wide Web Conference, Amsterdam.

Herrera, P, Serra, X. Peeters, G. (1999). Audio descriptors and descriptors schemes in the context of MPEG-7. Proceedings of the 1999 ICMC, Beijing, China, October.

Pachet, F. Roy, P. Cazaly. D. (2000) “A Combinatorial approach to content-based music selection”. IEEE Multimedia, pp. 44-51, March.

Pachet, F. and Laigre, D. (2001) A Naturalist Approach to Music File Name Analysis. in Proceedings of 2nd Int. Symp. on Music Information Retrieval.

Pachet, F. Westerman, G. Laigre, D. (2001) Musical Data Mining for Electronic Music Distribution., Proceedings of WedelMusic 2001, Firenze (It).

Pachet, F. and Zils, A. (2003) “EDS: Evolving High Level Music Descriptors”, Sony-CSL Tech. Report 2003-01-01.

Peeters, G. Rodet, X. (2002) Automatically selecting signal descriptors for sound classification. Proceedings of the 2002 ICMC, Goteborg (Sweden), September.

Scheirer, Eric D. (1998) “Tempo and beat analysis of acoustic musical signals”, J. Acoust. Soc. Am. (JASA) 103:1 (Jan 1998), pp 588-601.

Scheirer, Eric and Slaney, Malcolm. Construction and evaluation of a robust multifeature speech/music discriminator. Proc ICASSP ’97, pp. 1331-1334.

Wold, E. Blum, T. Keislar, D. Wheaton, J. (1996) Content-Based Classification, Search, and Retrieval of Audio, IEEE Multimedia, 3:3, pp. 27-36.

Zils A., Pachet F., Delerue O., Gouyon F. (2002) Automatic Extraction of Drum Tracks from Polyphonic Music Signals. Proc. of WEDEL MUSIC 2002, December.

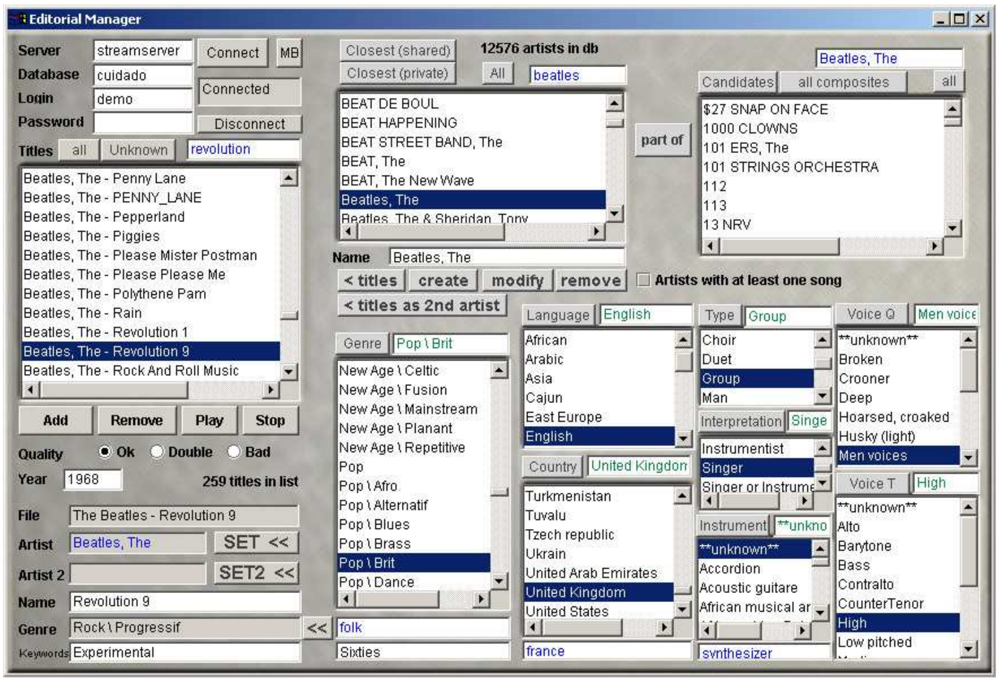

Figure 1. The Editorial Data Manager. The artist "Beatles" is selected, and corresponding metadata is shown (e.g. genre is "Pop / Brit"). The song "Revolution 9" is selected and its editorial information is shown as well (e.g. genre is "Rock / Progressive"). Keywords are associated both to the artist and song genres. Buttons allow to fetch artist and song information either locally (private data), or from the global Cuidado server (shared data).

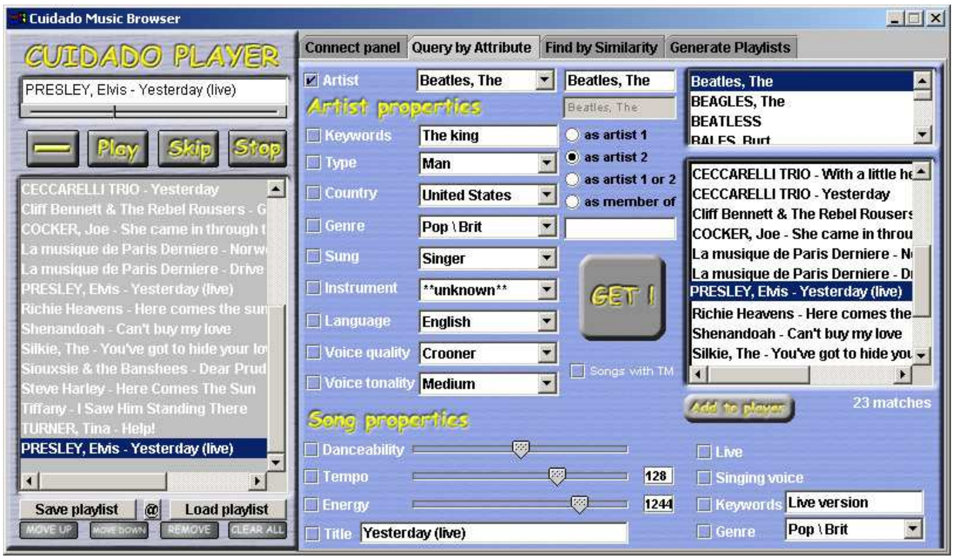

Figure 2. The Query Panel of the Music Browser. All the descriptors, editorial (e.g. genre) or acoustic (e.g. tempo, energy, danceability) can be used to issue queries. Here, the query is to get all the songs which have the Beatles as secondary artist (i.e. composer). Then the live cover of Yestesday by Elvis Presley is selected.

Figure 3. The Similarity Panel of the Music Browser. The similarity relation can be selected by the user (timbral, cultural from a web crawler, or inferred from radio playlists). The results are projected onto a similarity distance built from editorial information on artists and their genres.

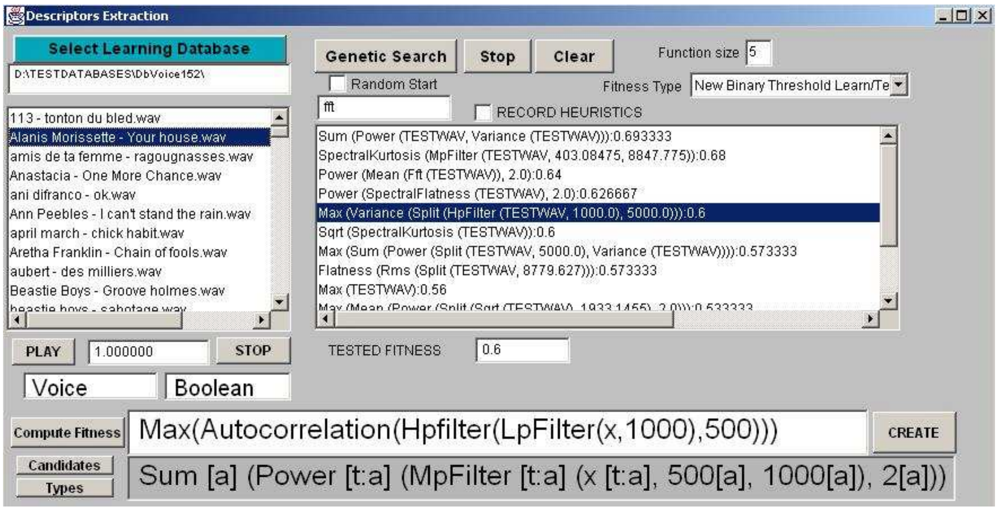

Figure 4. The EDS system finds automatically signal processing functions that extract high-level descriptors from acoustic signals. On the left part, a database of songs for which human perceptive tests have been performed is used to evaluate functions. EDS contains heuristics and rewriting rules to guide its search and simplify expressions. Here the expression in the text field is simplified, and its type is inferred (below).