

Good Vibrations: Music Discovery through Personal Musical Concepts

Vegard Sandvold, Thomas Aussenac, Òscar Celma, Perfecto Herrera

Universitat Pompeu Fabra

Pg. Circumval · lació 8

Barcelona, 08003 Spain

vsandvold, taussenac, ocelma, pherrera@iua.upf.es

Abstract

We present here Good Vibrations, a tool for music tagging, exploration and discovery, shaped as a media player plugin, and intended for home users. The plugin allows the quick "invention" of concepts and properties that can be tagged to songs. After some hours of active tagging, the plugin starts automatically proposing the proper tags to the user, who is also allowed to correct them. The plugin generates playlists according to the user-defined concepts, and recommends related music either from the user's personal collection or from the Internet (through its connection to Foafing the Music). The plugin runs, for the moment, in Nullsoft Winamp on Windows XP systems.

Keywords: Playlist generation, music recommendation, music tagging, software tools.

1. Introduction

Music taste is for many strongly connected to a sense of identity. This connection is reinforced whenever we buy a new album for our CD collection, participate in a music-related discussion forum, publish playlists on the Internet, or let software track our listening habits. Music creates communities and bonds between individuals with similar taste, and we rely on credible personalities and proven entities, *tastemakers*, to expose us to new and relevant music.

A number of music recommendation systems are available today – Last.fm, All Music Guides, MusicStrands, MoodLogic, MusicIP and Musicmobs, just to name a few. With a few exceptions, these systems are based on collaborative filtering, manual annotation or a combination of both. Foafing the Music[1] is an example of more novel approaches, where user profiles are combined with context-based web information and content-based audio descriptors.

In this demonstration paper we present Good Vibrations, a music exploration and discovery tool with powerful playlist generation capabilities. It is implemented as a pluggable program module (plugin) for media players (currently available for Winamp), operating mainly on personal music col-

lections. Good Vibrations draws its strength from a novel user interaction process, where the user tags songs with associations to semantic concepts such as mood, intensity or genre. The system quickly adapts to the user's style of tagging (or personal preference), and starts to automatically propose the proper tags.¹ The automatic tagging is done by predictive models, built from content-based audio descriptors and the annotations already provided by the user. Playlists are generated according to the user-defined concepts, either from the user's personal collection or from the Internet. Online recommendation is provided by Foafing the Music[1].

This conceptual categorization process can also be seen as a process of macro-annotation. By assigning semantic labels to songs on a global level, a data set is collected that in turn is used to train predictive models. Good Vibrations can thus be seen as a continuation of the MUCOSA Collection Tagger[2], intended for the home user.

2. Developing Musical Concepts

People develop personal musical concepts by making associations between songs and conceptual categories. Listening to "Like a Virgin" by Madonna may activate concept categories like up-beat, carefree, sexy, female vocalist, bass guitar etc. Some of these concepts are directly related to the content of the music, some to the artist or a particular listening experience, while others are related to the values a given social group assigns to them. Good Vibrations lets the user record associations to musical concepts, much like the conceptual categorization process of our cognitive system. The associations form the basis for predictive models used to organize, discover and recommend music.

The concept tagging dialog in Figure 1 offers the following functionality:

- Artist, album and title information for the currently playing song.
- Links to additional information on MP3.com²

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

© 2006 University of Victoria

¹ The particular brand of tagging used in Good Vibrations is not identical to free text tagging as we know it from flickr, del.icio.us and last.fm, but it brings many of the same positive effects (flexibility, playfulness and self-revelation) into the categorization process. Any number of concepts can be tagged to a song, and concepts are easily created or modified.

² <http://www.mp3.com/>

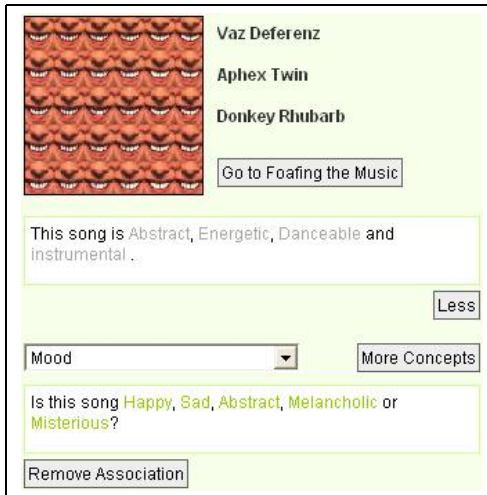


Figure 1. Good Vibrations concept tagging dialog.

- Album cover art downloaded from a web service provided by Foafing the Music[1].
- Associated concept categories, formatted like a text in natural language. For the Madonna example, the text may read “This song is up-beat, moderate and care-free, with a female vocalist.” The color of a category label indicates if the associations were tagged by the user or automatically by the system.
- Accessible tagging and concept manipulation operations (extend, modify and create new concepts).

In order to “jump-start” the tagging process, Good Vibrations will from the very beginning propose tags to the user, using a set of pre-defined predictive concept models built from a general collection of manually annotated music. In the beginning the user may disagree with automatically proposed tags, as conceptual categories often are idiosyncratic, with differences depending on a person’s individual history. As the user corrects proposed tags, the predictive models gradually converge towards the user’s true concept definitions, and the system gets increasingly better at tagging.

3. Creating Playlists

The aim of Good Vibrations is to provide playlists and recommendations based on music similarity. A satisfactory universal measure of similarity for music can be difficult to establish because the relevant variables for similarity will depend on a person’s conceptual categories. By using personal musical concepts as basis for music similarity computations, recommendations with a higher level of personal relevance can be provided.

All playlists can be saved to file, or loaded directly into the media player. The playlist entries display a song’s artist, album, title and length, and allows bookmarking for later retrieval, short previews and enqueueing and playing songs in the media player. Playlists can be generated according to the following modes and constraints:

- Songs similar to the currently playing song, sorted according to similarity, using a distance measure developed for the UPF MusicSurfer[4].
- Songs with similar tags to the currently playing song, or songs going from one category of a concept to another (e.g. from sad to happy). Since every song in the user’s collection can be the seed of a playlist, the potential for variation is proportional to the size of the collection.
- A special playlist listing all bookmarked songs.
- Recommendations for free online music, provided by Foafing the Music[1].

For maximum convenience, Good Vibrations scans the library of the media player for audio files, which are consequently analyzed and classified using the predictive concept models. Descriptors are extracted from the audio representing properties of rhythm, harmony, timbre and instrumentation, intensity, structure and complexity[3], running 15 times faster than real-time on a 2.4 GHz Pentium processor. Automatic tagging is performed by a classifier implementing nearest neighbor heuristics.

4. Pending work

Good Vibrations will soon be ported to other platforms and media players. The plugin is usable and well-functioning, but it remains to conduct a systematic evaluation of the usability and user satisfaction. We believe a web site for sharing personal musical concepts will be a meaningful supplement to playlist sharing. This “musical culture” site will allow users to upload, rate, comment and download each others predictive concept models.

5. Acknowledgments

The research and development reported here was partially funded by the EU-FP6-IST-507142 SIMAC (Semantic Interaction with Music Audio Content) project. The contributions made by Bram de Jong, Nicolas Wack, Amaury Hazan, Jose Pedro Garcia and Pedro Cano are highly appreciated.

References

- [1] Ò. Celma, M. Ramirez and P. Herrera, “Foafing the Music, A Music Recommendation System Based on RSS Feeds and User Preferences,” in *Proceedings of the 6th International Conference on Music Information Retrieval*, London, 2005.
- [2] P. Herrera et al. “MUCOSA: A Music Content Semantic Annotator,” in *Proceedings of the 6th International Conference on Music Information Retrieval*, London, 2005.
- [3] P. Herrera et al. “SIMAC: Semantic Interaction with Music Audio Content,” in *Proceedings of the 2nd European Workshop on the Integration of Knowledge, Semantic and Digital Media Technologies*, Savoy Place, London.
- [4] P. Cano et al. “Content-based Music Audio Recommendation,” in *Proceedings of ACM Multimedia 2005*, Singapore, Singapore.