SONGEXPLORER: A TABLETOP APPLICATION FOR EXPLORING LARGE COLLECTIONS OF SONGS

Carles F. Julià, Sergi Jordà
Music Technology Group
Universitat Pompeu Fabra, Barcelona, Spain
{carles.fernandez, sergi.jorda}@upf.edu

ABSTRACT

This paper presents SongExplorer, a system for the exploration of large music collections on tabletop interfaces. SongExplorer addresses the problem of finding new interesting songs on large music databases, from an interaction design perspective. Using high level descriptors of musical songs, SongExplorer creates a coherent 2D map based on similarity, in which neighboring songs tend to be more similar. All songs are represented as throbbing circles that highlight their more relevant high-level properties, and the resulting music map is browseable and zoomable by the users who can use their fingers as well as specially designed tangible pucks, for helping them to find interesting music, independently of their previous knowledge of the collection. SongExplorer also offers basic player capabilities, allowing the users to organize the songs they have just discovered into playlists which can be manipulated as well as played and displayed. In this paper, the system hardware, software and interaction design are explained, and the usability tests carried are presented. Finally, conclusions and future work are discussed.

1. INTRODUCTION

Since the popularization of the Internet and broadband connections, the amount of music which we are exposed to, has been increasing permanently. Nowadays, many websites do offer very large collections of music to the user, either free of charge (e.g. Magnatune [1], Jamendo [2]) or on a fee-paying basis (e.g. iTunes [3], The Orchard [4]). Such a number of available and still undiscovered music records and songs seems too difficult to manage in a sorting and searching-by-keyword way. In order to solve this problem and help users to discover new music, many online music recommendation services have been created (e.g. Pandora [5], Last.fm [6]). One of the main drawbacks of most current music recommenders, independently of the recommendation mechanisms and algorithms they employ (user profiling, experts-based knowledge, statistical models, etc.), is that they apply information filtering techniques to the entire collections, in order to extract and display only a subset of songs that the system believes the user could enjoy. By doing it this way, the user loses the opportunity to discover many new songs which are not presented by the system, whatever the cause may be.

To solve this problem, we propose to construct maps of the entire collections of songs and allow users to explore them in novel ways. Maps are widely used to explore spaces and also concepts. Although most commonly used to depict geography, maps may represent any space, real or imagined, without regard to context or scale. We use conceptual maps to discuss ideas, we organize data in 2D spaces in order to understand it, and we can get our bearings using topographical maps. SongExplorer's maps are constructed using MIR techniques that provide the high-level descriptors needed successfully organizing the data; they do not filter or hide any content, thus showing the complete collection while highlighting some of the songs' characteristics.

Therefore, SongExplorer provides intuitive and fast ways for promoting the direct exploration of these maps. In the last years, several successful projects have shown that tangible, tabletop and multitouch interfaces exhibit useful properties for advanced control in general (such as continuous, real-time interaction with multidimensional data, and support for complex, skilled, expressive and explorative interaction) [4] and for the exploration of bidimensional spaces in particular [2]. Following this trend, SongExplorer allow users to interact with the maps directly with their hands, touching the surface with their fingers and manipulating physical tokens on top of it. In the following section we will comment some of the most relevant previous works, related to the two main aspects of our project, i.e. (i) the visualization of musical data, and (ii) the direct manipulation of this or any other type of data, in a tabletop interaction context.

[1] http://www.margatune.com
[2] http://www.jamendo.com
[3] http://www.apple.com/itunes/
[4] http://www.theorchard.com
[5] http://www.pandora.com
[6] http://www.last.fm

2. RELATED WORK

2.1 Visualization of music collections

In the field of visualization, there is an extensive bibliography on the representation of auditory data. In the particular case we are focusing on, that of the visual organization of musical data, solutions often consist in extracting feature descriptors from data files, and creating a multidimensional feature space that will be projected into a 2D surface, using dimensionality reduction techniques.

A very well known example of this method is the work Islands of Music by Pampalk [13], which uses a landscape metaphor to present a large collection of musical files. In this work, Pampalk uses a Self Organizing Map (SOM) [9] to create a relief map in which the accumulation of songs are presented as the elevation of the terrain over the sea. The islands created as a result of this process roughly correspond to musical genres.

A later attempt to combine different visualizations on a single map was also created by Pampalk et al [14]. By using different parameters to organize the SOM, they created several views of the collection, later interpolating the different solutions for creating a smooth combination of situations with which to explore new information.

Beyond the 2D views, an interesting work on music collections visualization, which distributes the songs on a spherical surface, thus avoiding any edge or discontinuity, is described by Leitich and Topf [11].

In the aforementioned examples, a topological metaphor is taken in advantage to enable users exploring big collections of data. A different and original visualization approach is chosen in Musicream [1], an interesting example of exploratory search in music databases, using the search by example paradigm. In Musicream, songs are represented using colored circles, which fall down from the top of the screen. When selected, these songs show their title on their center, and they can be later used to "fish" similar ones.

2.2 Tangible tabletop interaction

In the domain of Tabletop and Tangible User Interfaces (TUI) there is also a growing interest in working with musical applications. As a matter of fact, from the Audiopad [15] to the Reactable [5], music performance and creation has arguably become the most popular and successful application field in the entire lifetime of this interaction paradigm. This is, according to Jordà [4], because of the specific affordances of this type of interfaces: support of collaboration and sharing of control; continuous, real-time interaction with multidimensional data; and support of complex, expressive and explorative interaction. In this sense, and although less prolific than the applications strictly conceived for musical performance, some interesting works have also been developed to interact with large music collections.

Musictable [16] takes a visualization approach similar to the one chosen in Palmpalk's Islands of Music, to create a two dimensional map that, when projected on a table, is used to make collaborative decisions to generate playlists. Another adaptation into the tabletop domain is the work from Hitchner et al [3], which uses a SOM to build the map and also creates a low resolution mosaic that is shown to the user. The users can redistribute the songs on this mosaic and thus changing the whole distribution of SOM according to the user's desires.

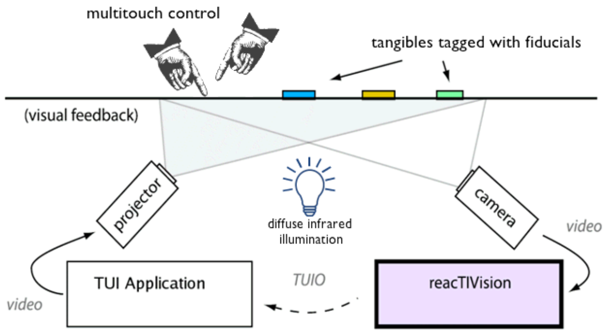

Figure 1. reactiVision framework schema.

We believe this paper also represents a real contribution to the tangible/tabletop user interface community. As noted before, it has been proposed very recently [4] that they can be especially adequate for complex and non-task oriented types of interaction, which could include real-time performance, as well as explorative search. The topic addressed by this paper (N-Dimensional navigation in 2-D) has never been addressed before within tabletop interfaces.

3. HARDWARE

SongExplorer is a tabletop application, i.e. a computer application meant to run on a tangible and multitouch surface, designed for the exploration and discovery of new music. In this section we discuss its main hardware components.

As schematized in Fig.1, the system is made of a translucent plastic surface, some infrared lamps for diffused illumination, an infrared camera for the detection of the user interaction, and a projector for the projection of the visual feedback on the table surface. The surface is round, as in the Reactable case, for encouraging collaboration [5].

The tracking software is based on reacTIVision [6], an open-source framework for the recognition of fingers and objects tagged with fiducials. The images showing the fiducial markers that are stuck into the physical pucks, and the fingers that are in contact with the translucent surface, are captured by the infrared camera and processed by reacTIVision. For each video frame, this software component sends the corresponding data (which includes the positions and IDs of the identified objects and fingers) to SongExplorer, using the TUIO protocol [7]. SongExplorer subsequently identifies the gestures and the actions performed on the table surface, and proceeds with the appropriate responses, finally generating the output image that is displayed by the projector on the translucent surface.

4. SOFTWARE

This section describes the main components of the SongExplorer software: feature extraction, visualization and interaction.

4.1 Feature Extraction

SongExplorer uses all the songs included in the Magnatune online database, which comprises a total of 6666 songs weighting more than 26 GB. Being Creative Commons-licensed, this library is used in many research projects. These songs are processed by an in-house music annotation library developed at the Music Technology Group (MTG) [10], and the results are transformed to binary files that can be loaded by the system using the Boost[7] C++ library.

4.2 Visualization

From the whole set of available annotated features generated by the annotation library, we are curently using the most high-level properties together with the BPM:

- Beats Per Minute (BPM)
- Happy probability
- Sad probability
- Party probability
- Acoustic probability
- Aggressive probability
- Relaxed probability

All these high level features are independent, and even the moods, which try to cover all the basic emotions, do not depend on each other (i.e. a song could be both sad and happy) [10]. The emotional features can, in fact, be considered binary, with their values indicating the probability of this feature being true.

With this data, a multidimensional feature space (of 7 dimensions) is constructed, in which each song is a single data point with its position defined by these 7 features, all of them being normalized between 0 and 1. From this multidimensional data we construct a 2D space which preserves its topology, and we present it to the user, who will then be able to explore it.

Similarly to other visualization works, a SOM is used to distribute the data on the 2D space. Our implementation of the Kohonen network uses a circular, hexagonally-connected neuron grid, in order to fit the shape of the interactive surface. As opposed to the original implementation of SOM [9], a restriction was added to prevent more than one song falling into a single neuron, so that every representation in the 2D space should be visible and equally distant from its direct neighbors, as shown in Fig. 2.

In the visualization plane, every song is represented by a colored circle, throbbing at the song's BPM. Since there seems to be a strong agreement about the usefulness of artwork to recognize albums or songs [11, 12], depending on the zoom factor, the actual artwork may be shown in the center of each song.

[7] http://www.boost.org

Figure 2. Detail of the hexagonal structure of the grid.

Additionally, colors are used to highlight the different properties of the songs. The coupling {feature → color} was defined with an online survey where 25 people had to pair the high-level tags to colors. The color candidates were 7 basic colors with maximum saturation and lightness: red, blue, green, brown, cyan, yellow and magenta. Subjects were only able to choose the best color representation for each tag. The results were: aggressive-red (with an agreement of 100%), relaxed-cyan (43.5%), acoustic-brown (52%), happy-yellow (39%), party-magenta (48%) and sad-blue (56.5%).

For every song, its corresponding property value is mapped into the saturation of the related color (0 meaning no saturation thus resulting on a grey color, 1 corresponding to full saturation), while the lightness is kept to the maximum and the hue is obviously linked to the emotional feature se-

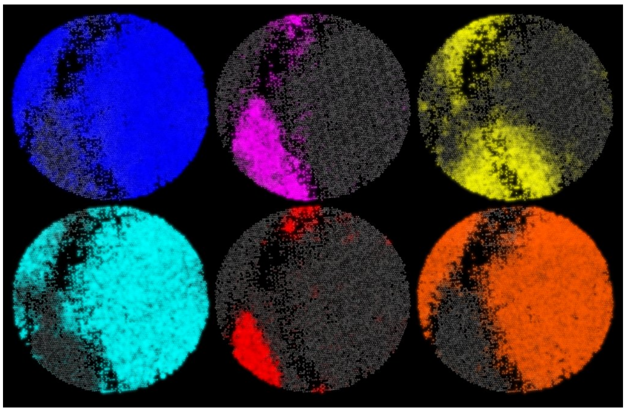

Figure 3. Colors highlighting high-level properties: sad, party, happy, relaxed, aggressive and acoustic (Best seen in color, colors modified for B/W printing).

Figure 4. The tangibles of SongExplorer: *playlist navigator*, *color changer*, *magnifying glass* and *navigation menu*

lected, as described in the previous color pairings (Fig. 3 shows the effect of different highlights on the songs). An option to see colors representing genres is also provided, although in that case the pairing between genres and colors is done randomly.

4.3 Interaction

From a users perspective, SongExplorer is a table that shows dynamic images on its surface, which can be manipulated in several ways, using both the fingers as well as some special pucks we will call tangibles, and which will be described later.

4.3.1 Multitouch interaction

Basic finger interaction includes single and multiple finger gestures, and the use of one or two hands. The simplest gesture, selecting and clicking, is implemented by touching a virtual object shown on the table surface, with a single finger and for more than 0.4 seconds. In order to distinguish them from the selection action, other finger gestures involve the use of two simultaneous fingers for each hand. That way, using only one hand, users can translate the map and navigate through it, while the use of both hands allows rotating and zooming the map (see Fig. 5). It should be noted that most of these gestures have become de-facto standards in multitouch and tabletop interaction [8].

4.3.2 Tangible interaction with pucks

Additionally, SongExplorer tangibles also include 4 transparent Plexiglas objects of about 50cm2 each, each one with a different shape and a different icon that suggests its functionality, as described in Table 1. These pucks, which can be kept on the table frame outside the interactive zone (see Fig. 4), become active and illuminated when they get in contact with the interactive surface. As indicated below, some (like the color changer or the navigator) will apply to the whole map, while others (such as the magnifying glass) apply to the selected song.

- The **color changer** puck allows selecting and highlighting one of the different emotional properties of the whole song space. For example, changing the map to red allows us to see the whole map according to its aggressive property, with fully red dots or circles corresponding to the more aggressive songs, and grey dots to the least aggressive ones. Apart from helping to find songs based on a given property, the resulting visual cues also help to memorize and recognize the explored zones of the map.

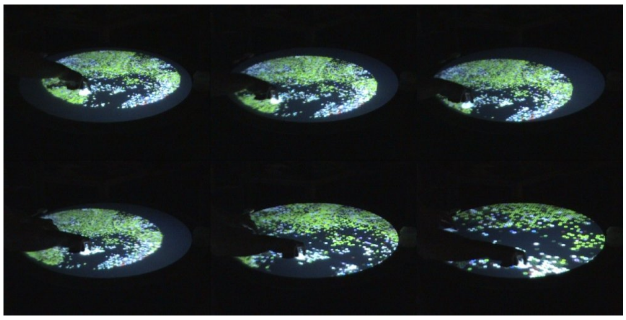

Figure 5. Virtual Map movement (up) and zooming (down)

- When placed on top o a song, the **magnifying glass** puck allows seeing textual information on this particular song, such as the song title, the album, the authors name, as well as the artwork.
- The **navigation** puck displays a navigation menu, which allows the user to perform actions related to the movement and zooming of the map, such as returning to the initial view, or zooming and centering on the currently playing song.
- The **playlist navigator** puck allows the creation and management of the playlist, as described below.

4.3.3 Managing playlist and playing back songs

SongExplorer has the ability of creating and managing playlists. Playlist are graphically represented on the surface as a constellation, in which the stars (i.e. the corresponding songs it contains) are connected by lines establishing their playing order (see Fig. 6). Most stars show a white stroke, except for the one that is currently playing (red), and the one the playlist navigator is focusing on (green).

Playlists allow several actions using both the fingers and the playlist navigator puck. When clicking on a song, this is automatically added to the playlist. Users can start playing a song by clicking on any star of the playlist. Similarly, crossing out a star removes the corresponding song from the list. A song will stop playing either when it reaches its end, when the song is deleted from the playlist or when another song is selected for playing, and a playlist will keep playing until its end, unless it is stopped with the playlist navigator puck. This object allows several additional actions to be taken on the playlist, such as navigating through its songs and showing information about the focused song in the same way the magnifying glass does.

5. EVALUATION

Some user tests have been undertaken in order to evaluate the system, focusing on the interface design. The evaluation consisted in three areas: subjective experience, adequacy of the visualization and the organization, and interaction.

Symbol	Name	Description
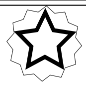	playlist navigator	Permits to run over the songs on the playlist
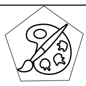	color changer	Allows to highlight features of the songs
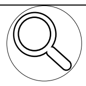	magnifying glass	Shows information about songs
	navigation menu	Provides a way to return to known situations

Table 1. Tangibles used in SongExplorer

Figure 6. Playlist and Playlist navigator

5.1 Experiment Procedure

To carry out the tests, an interactive table with SongExplorer up and running was provided. The system was always on an initial state at the beginning. One subject at a time was doing the test. First of all, a little explanation about the purpose, visualization and interaction was given. Then the subject was asked to *find something interesting* in the music collection. No time limit was imposed, and the subject was observed along the process. At the end of the activity, the subject was told to fill a questionnaire, on which she had to rate, using a Likert scale of 10 levels [8], the several aspects of each area. They could also write suggestions at the end of the test.

5.2 Results

After doing the tests the results were quite positive (see Table 2). Regarding the personal experience with SongExplorer, the subjects enjoyed the experience, discovered new and interesting music, felt comfortable, and found it useful to find interesting music. So the overall experience seemed to be good; we have to notice the low deviation, indicating that there was an agreement about these opinions.

	$\mu_{1/2}$	IQR
Enjoyed the experience	8	1
Discovered new music	8	1
Felt comfortable	8	1.5
Found it useful	9	0.5
Found colors correct	8	1.5
Found categories suitable	7	1
Found graphics adequate	9	1.5

Table 2. Evaluation Results. $\mu_{1/2}$: Median, IQR: Interquartile range.

Focusing on the visualization process, there was also a common opinion about the suitability of the colors used. This is not a surprise, as they were extracted from an online poll (details on subsection 4.2). The categories (formerly the high-level properties from the annotation software) were suitable, according to the subjects, for the purpose of describing music. The graphics were also evaluated (meaning the adequacy of icons, the metaphor songcircle, the panels...) and also appreciated.

For the evaluation of the interaction, this paper will not enter into details, because of its extension, but the results were also quite positive. The level of understanding of every gesture and tangible of SongExplorer was tested, as well as their difficulty of use and usefulness. The only noticeable result was that there seemed to be an inverse correlation between previous experiences with tabletops and the perceived difficulty of finger gestures.

Finally there was a general demand for more music-player capabilities like pause or a progress bar for jumping to the middle of the song. The option of bookmarking and storing playlist was also desired.

6. CONCLUSIONS AND FURTHER WORK

We have presented SongExplorer, a new system for large music collections exploration, based on similarity and high level property highlighting that can allow users to find interesting new music.

[8] 10: Totally agree, 0: Totally disagree.

The user tests have shown that this system can be a good tool for discovering new, valuable music to the users. And this forces us to think about its possible real world applications. As long as this type of interfaces are uncommon, it is not intended for personal use because of its physical nature (size) and its hardware requirements. But other uses than the personal one can be imagined. For instance, some researchers from the annotation software communicated their desire to use SongExplorer to test the reliability of its annotation systems. Using the virtual map they can easily search for inconsistencies. This can be extended to other annotation software systems.

As another real world user case, it would be useful, as a way of promoting music in stores, to have this system available to their customers. An additional feature could be created allowing users to highlight their favorite music so they can then find similar music near to the ones they like, to optionally buy the records afterwards.

In the future versions of SongExplorer, we want to give it the ability of storing playlists, give the user the option of rating songs, adding common player-like capabilities like jumping to the middle of a song, searching songs using actual records (using identifiers on the CD cases) and probably more features.

Video:

`http://www.vimeo.com/4796964`

7. ACKNOWLEDGEMENTS

We thank the MTG team at Universitat Pompeu Fabra for its excellent work on music mood classification that made this project possible. Also to the Reactable project for providing access to its hardware and software platform that is the base of SongExplorer. Finally, we want to thank the open source community for providing the software libraries and free content (Magnatune) on witch this application relies.

8. REFERENCES

[1] M. Goto and T. Goto. Musicream: New music playback interface for streaming, sticking, sorting, and recalling musical pieces. In *ISMIR 2005: Proceedings of the 6th International Conference on Music Information Retrieval*, 2005.

[2] J.Y. Han. Multi-touch interaction wall. In *International Conference on Computer Graphics and Interactive Techniques*. ACM New York, NY, USA, 2006.

[3] S. Hitchner, J. Murdoch, and Tzanetakis G. Music browsing using a tabletop display.

[4] S. Jordà. On stage: the reactable and other musical tangibles go real. *International Journal of Arts and Technology*, 1(3):268–287, 2008.

[5] S. Jorda, M. Kaltenbrunner, G. Geiger, and R. Bencina. The reactable*. In *Proceedings of the International Computer Music Conference (ICMC 2005), Barcelona, Spain*, pages 579–582, 2005.

[6] M. Kaltenbrunner and R. Bencina. reacTIVision: a computer-vision framework for table-based tangible interaction. In *Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 69–74. ACM New York, NY, USA, 2007.

[7] M. Kaltenbrunner, T. Bovermann, R. Bencina, and E. Costanza. TUIO: A protocol for table-top tangible user interfaces. *Proc. of the The 6th Int'l Workshop on Gesture in Human-Computer Interaction and Simulation*, 2005.

[8] J. Kim, J. Park, H.K. Kim, and C. Lee. Hci (human computer interaction) using multi-touch tabletop display. In *IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, 2007. PacRim 2007*, pages 391–394, 2007.

[9] T. Kohonen. *Self-Organizing Maps*. Springer, 2001.

[10] C. Laurier, O. Meyers, J. Serrà, M. Blech, and P. Herrera. Music mood annotator design and integration. Chania, Crete, Greece, 03/06/2009 2009.

[11] S. Leitich and M. Topf. Globe of Music: Music Library Visualization Using GEOSOM. In *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007).*

[12] A. Pabst and R. Walk. Augmenting a rugged standard DJ turntable with a tangible interface for music browsing and playback manipulation. In *Intelligent Environments, 2007. IE 07. 3rd IET International Conference on*, pages 533–535, 2007.

[13] E. Pampalk. Islands of Music Analysis, Organization, and Visualization of Music Archives. *Journal of the Austrian Soc. for Artificial Intelligence*, 22(4):20–23, 2003.

[14] E. Pampalk, S. Dixon, and G. Widmer. Exploring music collections by browsing different views. *Computer Music Journal*, 28(2):49–62, 2004.

[15] J. Patten, B. Recht, and H. Ishii. Audiopad: A tag-based interface for musical performance. In *Proceedings of the 2002 conference on New interfaces for musical expression*, pages 1–6. National University of Singapore Singapore, Singapore, 2002.

[16] I. Stavness, J. Gluck, L. Vilhan, and S. Fels. The MUSICtable: A Map-Based Ubiquitous System for Social Interaction with a Digital Music Collection. *Lecture Notes In Computer Science*, 3711:291, 2005.