Music Listening History Explorer

An alternative approach for browsing Music Listening History Habits

Ricardo Dias, Manuel J. Fonseca, Daniel Gonçalves

Department of Information Systems and Computer Science

INESC-ID/IST/Technical University of Lisbon

Lisbon, Portugal

ricardo.dias@ist.utl.pt, mjf@inesc-id.pt, daniel.goncalves@inesc-id.pt

ABSTRACT

Nowadays, people spend time using services to track their music listening history. Although these services provide statistics and small graphics/charts, they are mainly used to record and to allow direct access to the information, not providing any visualization and exploration functionality. In this paper we describe a new approach for browsing and visualizing music listening histories, which combines a timeline-based visualization, with a set of synchronized-views and an interactive filtering mechanism to provide a flexible and easy to use solution. This was complemented with brushing and highlighting techniques that allow users to observe trends on artists, albums and tracks listening. Experimental evaluation with users revealed that they were able to complete all the proposed tasks with a low error rate, and that they found the solution easy to use. Moreover, users liked our approach for browsing and exploring listening histories, emphasizing its flexibility and effectiveness, and founding the full experience engaging and rewarding.

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: Graphical User Interfaces - GUI

General Terms

Design, Experimentation

Author Keywords

Listening History, Interactive Browsing, Visualization

INTRODUCTION

Nowadays, millions of people spend their time using online services to track different aspects of their lives. Moods[1], the places they go[2] or the songs they ear[3], are just a few examples. This data contains intrinsic information about users'

---

[1] http://www.moodstats.com
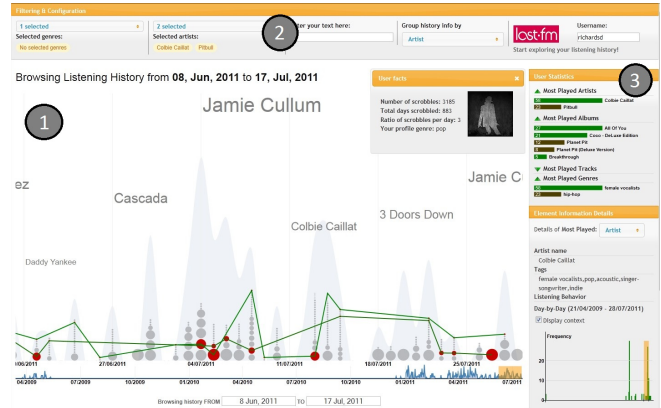
[2] https://foursquare.com

[3] http://www.last.fm

Figure 1: Prototype user interface.

habits and behaviors, and most of it can be tracked without user intervention. As a consequence, it quickly becomes quite complex to understand and visualize this type of personal information.

These services often enrich the available data with statistics and small graphics, but their main use is still to record and allow direct access to the information. This raises the opportunity to create visualizations using this type of personal information. Some visualization techniques and solutions have already been developed, being the most relevant those developed by Byron and Wattenberg, the *StreamGraph* [4], and those created by Baur et al, *LastHistory*[3], and the *Tangle, Strings and Knots*[2] visualizations. Other community-created static visualizations, such as the *Scrobbling Timeline*[4], the *Last.fm Explorer*[5] and the *LastGraph*[6], represent the fans effort to visualize this information, besides creating entertainment visualizations. However, these approaches present some limitations, such as: they are more concerned about creating general overviews and statically pleasing visualizations of listening histories, instead of providing interactive browsing and filtering mechanism; input data limitation, regarding the number of songs listened to; and finally, some rely on external software plugins or are OS dependent.

In this paper we describe a new approach for interactively

---

[4] http://playground.last.fm/demo/timeline

[5] http://alex.turnlav.net/last_fm_explorer/
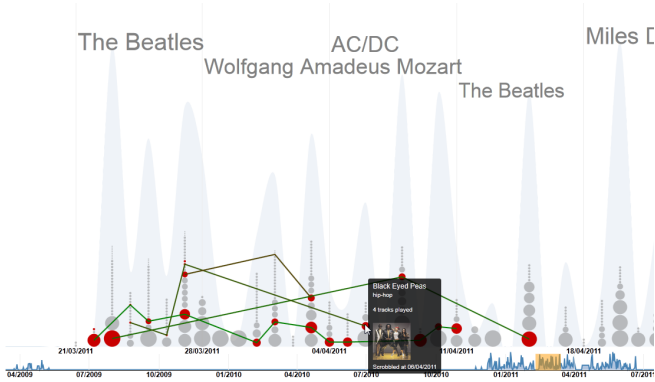
[6] http://lastgraph.aeracode.org

Figure 2: Multiple selection in main visualization.

browsing and filtering music listening histories that addresses the described concerns. Our solution uses data collected from the *Last.fm* service to create a compound visualization using a zoomable timeline and a set of synchronized views (see Figure 1). This solution provides not only an overview of the entire music listening history, but also an interactive mechanism to inspect a specific time period or *data element* (artist, album, track). It also supports listening histories up to thousands of tracks, and was built as a web application, so it could be available from anywhere.

Experimental evaluation shows that users were able to complete all the tasks with a very low error rate (<3%), and that they enjoyed using our solution and found it very easy to use. They also pointed that the solution was very flexible providing an engaging and rewarding experience. Results also revealed the users' awareness for this type of personal information, and the need to make use of it to personalize and improve music recommendation algorithms. The interactive exploration and filtering mechanisms were considered as the greatest advantage of our approach.

In the remainder of this document, we start by describing our solution and its design rational, providing insight about its main features and interaction. Later, we describe the conducted evaluation and present the main results. Finally, we conclude with our contributions and limitations, and present some directions for future work.

OUR SOLUTION
Our solution is a novel approach for exploring and visualizing music listening histories. On a timeline-based visualization users can browse their listening history in three different ways: *artists, albums and tracks*. Listening data can also be grouped by time and be filtered using the available criteria.

The developed prototype is shown in Figure 1. It includes three areas of interaction: main visualization area (1), filtering area (2) and statistics and details area (3).

Main Visualization
The visualization area (1) contains the *timeline* and the *stacked dot* visualizations for exploring the listening histories. As listening history data is inherently time based, tem-

poral and timeline-based techniques were chosen to visualize this information (see [1] and [5] for a context).

The *doted* visualization is the main visualization over the users' listening histories. In this technique, much alike the ones Viégas developed in [6] and [7], each bar of *stacked dots* encode a time slice of the selected time period, and each dot represents a *data element* (artist, album or track) listened in that time slice. For example, in Figure 1, each column represents a day, and each dot an artist listened in that day.

Listening history data can be grouped by *day, week, month or year*. Depending on the amount of data contained in the selected period, the visualization will change accordingly. As the user starts to increase the selected time period, the amount of data to be displayed in the same space also increases, making the size of the dots smaller. However, as this cannot continue until the visualization becomes imperceptible, whenever the dots maximum size is about to decrease above a determined minimum size, the group operation is changed to one higher level in the temporal hierarchy. In the given example, it would change from day to week, making the visualization perceptible again. This feature enables users to interactively browse the listening history by time, always using the best suitable visualization over data. Nevertheless, users always have control over the applied zoom and can switch between group modes whenever they need more precise details.

The size of the dots and their order are mapped by a function that uses the listening frequency of the encoded *data element*, and its relevance in the full listening history. The color of the dot, unlike what other solutions do, instead of representing the genre of the element, introduces a new concept to characterize the users: it represents how recent or how old the tracks listened are (regarding to the release date of the track/album). The colors range from *strong and warm* ones (e.g. red tones), representing recently release tracks/albums, to colder colors (e.g. blue tones), representing older elements. This property allows users to get insights about their listening habits, for example, "*Do I hear more recent musics than old ones?*", or "*Last month I heard old tracks from Bon Jovi, but now I move forward to listen just the new albums! Didn't realize that!*", something that cannot be answered using current state-of-the-art solutions.

As for interaction, whenever the user *highlights* a dot, additional information is displayed, and a line appears, connecting all the dots that represent the same *data element*. This property makes it possible to visually follow the evolution of that element's listening. By clicking over it (or pressing the *CTRL* key for multiple selection, see Figure 2), a filter is applied using the dot's data, and the visualization dynamically adapts to show the filtered information. The line colors are mapped by the ones used in the *user statistics* details, explained bellow.

Filtering and Browsing
The filtering area (2), contains the interface for the interactive filtering mechanism. Users can filter the main visualization using three different filters: *genre*, *artist name* and *free text*. These filters can be combined and results are *dynamically highlighted* in the visualization. Although the filters

limit the information we will see, they have different behaviors over the main visualization.

Genre filter makes the dots in the visualization to fade if they do not encode data from one of the selected genres, or to maintain its color otherwise. Although using this method users lose the ability to clearly identify the genre of each element, it gives them the ability to infer about the listened tracks/artists in that period and for that genre: "Were they recent or old?". When only one genre is selected, this fact becomes even clear, giving users the possibility to track their behavior for that genre over time.

As for artists filter, picking them from the set of listened artists in the selected time period, makes the respective dots to be get connected in the main visualization. Connecting the dots in the visualization acts as a visual clue to help users to track the trend of a specified artist, album or track listening (when in track/album visualization, artist filtering makes all the albums/tracks in the visualization to be individually connected, so users can follow the progress of a specific album/track listening). Although an increase in the line slope mean a decrease in frequency, during evaluation this did not prove to be a problem.

Text filter is a regular text input control where users can enter free text to filter data. It searches in the available tags and element information for matches (or partial matches), and then filters the visualization much alike the two filters above describe.

Statistics and Details

To provide an interactive environment to explore and filter users listening histories, our solution uses a synchronized-views paradigm. Besides the main visualization, two other provide statistic information and details about user listening history: the first one contains information about the most frequently played artists, albums, tracks and genres; and the other provides details about the most frequently played elements.

Statistics region displays general statistics about the most played artists, albums, tracks and genres, in the selected time period. By default, users can always keep track of the top five elements; however, by filtering the main visualization, they can obtain detailed information for less or more elements than those in the initial value. This visualization uses a *bar chart* encoded with the frequency an element was played, and a color mapping, from *green to dark red* (not collapsing with main visualization schema color), to indicate the most played element, and allow users to visually link with the other panels. These colors also map to the *brushing lines* that connect the dots in the main visualization.

In the details panel, users can obtain generic information, such as the release date of an album or the duration of the track, but also aggregated information specific to the user and the selected element. Two visualization techniques are used: a *day-by-day* and a hourly visualization (see figure 3). The day-by-day visualization consists in a *bar chart* where each bar represents a day in the selected time period when the element was listened to (see figure 3a), and the hourly visual-

(a) Day-by-day with no context

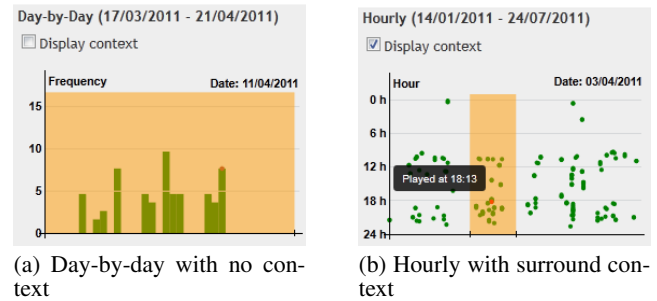(b) Hourly with surround context

Figure 3: Day-by-day and hourly visualizations.

ization in a *dot chart*, where each dot represents an element listened at a specific time (day, hour-minutes). For these two visualizations, the overall context can be displayed, see figure 3b, showing the total listening history for that element and highlighting the selected time period. These visualizations provide insights about how the user listened to that particular artist, album or track, answering questions like "*Which days did I play Colbie's music (and how much) and at what time? Do I hear it a lot? Always at morning?*".

EXPERIMENTAL EVALUATION

In order to evaluate our approach we conducted a set of user tests. We intended to evaluate the users' ability to complete all the proposed tasks, the error rate, and finally their satisfaction and experience about using our solution.

Users

We conducted tests with 10 users (2 females and 8 males), 7 of them aged between 20 and 30 years, and the remaining between 30 and 50 years. 7 out of 10 listen to music every day and have different backgrounds: four were graduate computer science students, two were software engineers (one being an amateur musician), one was a journalist, another was a human science student, and the remaining two were undergraduate computer science students (one worked as a saleslady in a local supermarket). 9 users reported listening to music at home, 8 at work and while driving. They all listen to music in computer and 8 of them using radios, and 7 use *smartphones* when moving around.

Procedure

Each test session lasted in average 30 minutes, and was performed using users' personal computers. The only requirement about the setup was that an Internet connection was available. To establish a baseline for all tests, a common *Last.fm* account was used for all sessions. Although this data was not personal to users, we adopted this solution because in these tests we do not intend to evaluate personal issues.

First, we started with a brief description about the test, thanking the person and explaining our objectives and what we were going to do. Next we introduced our application, explaining its main features and how to use them. At this time we gave a practice time for users to get familiar with the solution. A video was used in this part to present the interaction

and main features of the solution. Next, users performed a set of tasks using the application. Tasks covered different objectives as explained bellow. When authorized, user's interaction with the application was recorded, using a screen capture software for later post-analysis. Finally, users answered a satisfaction survey and an informal interview was conducted to get feedback from users about the experience.

Tasks and Objectives
Based on the analysis of related work (see [3]), we designed 10 tasks, grouped into two sets accordingly to our main objectives. In the first group, we tried to evaluate if users could use our approach to effectively explore the listening history. Tasks ranged from direct use of the applications interface (e.g., getting information about most played artists, or the period with more intense music listening activity), to others where combining filtering (with genre and artist information) was mandatory. On the second group, we intended to assess if users could find and use information that would lead them to get insights about *listening trends*, and to understand *statistics* about the user's profile. Tasks covered describing the listening evolution of a song (e.g., if it increased/decreased or was constant, etc.), or trying to understand some facets of the user's profile, such as, "*Does he listens to complete albums, or he prefers to hear isolated tracks?*".

Main Results
In short, all users were able to finish the required tasks in the given time period, and they did not found any relevant difficulty in executing them.

On average, 9 out of 10 users found the execution of the required tasks *very easy* (using a 5-point Likert scale ranging from *Very Difficult* to *Very Easy*), and only user 1, in task 5 about browsing found it to be tricky. This issue was due to a temporary confusion with the brushing lines concept.

Task success rate was very high for 9 of the 10 tasks. A task was considered successfully completed if the user finished it and gave a correct answer (when required). Indeed, only task 3 got an error rate of 40% against the average error rate of 3% for the remaining tasks. To complete this task, users should infer about the most intense listening period. This suggests that improvements should be made to provide more detailed information when necessary.

Satisfaction survey revealed some interesting results. 9 out of 10 users were able to use the application and understand the information presented. 4 of them had difficulties in executing task 7 (describing the listening history evolution of a song), but their main concern was about giving a wrong answer, as this was a subjective and interpretative question. When asked to classify the experience and the evaluation using semantic scales, 8 users described it as engaging and rewarding. They all felt that they were learning something about the other user's listening history, and wanted to use the application to explore their own listening history.

All users said that it was important to understand the listening history, but only three have previously tried it. They tried to obtain global statistics about the most played tracks/artists, and only one failed to get the desired information. They all said it would be easy to obtain it using our solution. When inquired about using this information to improve recommendation systems, for example, to suggest songs or generate playlists, they all agreed that this is essential, as this data encodes their tastes and listening patterns.

In general, users enjoyed using our solution and they intended to use it to explore their own listening history. One user even said: "*This tool is very nice. I would like to use it to browse my listening history. I wonder which my personal patterns are and how my behavior changes over time!*".

CONCLUSIONS AND FUTURE WORK
In this paper we described an alternative solution for exploring and filtering music listening histories. The solution used a timeline-based visualization and a set of synchronized-views to allow users to perform an interactive exploration of their music listening histories. Results from user evaluation showed that the developed solution was easy to use, with users completing all the tasks with a low error rate.

As for future work, we are willing to follow two paths: the first one is to conduct a further user evaluation to assess if users can get insights about their personal patterns and habits, while browsing and filtering their personal listening histories; next, we intend to apply data mining techniques to this information to create a set of user profiles (for example, the users that always seek the tops, the ones that enjoy mostly female voices or that listen to classic music in the morning, etc.), and then use them to improve and personalize music recommendation and automatic playlist generators.

REFERENCES
1. Aigner, W., Miksch, S., Müller, W., Schumann, H., and Tominski, C. Visualizing time-oriented data - a systematic view. *C&G* (June 2007).

2. Baur, D., and Butz, A. Pulling strings from a tangle: visualizing a personal music listening history. IUI (2009).

3. Baur, D., Seiffert, F., Sedlmair, M., and Boring, S. The streams of our lives: Visualizing listening histories in context. *IEEE TVCG* (2010).

4. Byron, L., and Wattenberg, M. Stacked graphs geometry & aesthetics. *IEEE TVCG* (2008).

5. Muller, W., and Schumann, H. Visualization methods for time-dependent data - an overview. In *Simulation Conference* (2003).

6. Viegas, F., Boyd, D., Nguyen, D., Potter, J., and Donath, J. Digital artifacts for remembering and storytelling: posthistory and social network fragments. In *HICSS '04* (2004).

7. Viégas, F. B., Golder, S., and Donath, J. Visualizing email content: portraying relationships from conversational histories. CHI (2006).