

# Interactive Exploration of Music Listening Histories

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

Over the past years, music listening histories have become easily accessible due to the expansion of online lifelogging services. These histories represent the sequence of songs listened by users over time. Although this data contains intrinsic users' tastes and listening behaviors, it has been mainly used to personalize recommendations. Tools to help users exploring and reasoning about the information contained in the listening history, only recently have started to emerge. In this paper we describe a new visualization and exploration tool that allows users to interactively browse their listening histories, while leading them to identify listening trends and habits. Our solution combines a rich-featured timeline-based visualization, a set of synchronized-views and an interactive filtering mechanism to provide a flexible, effective and easy to use system for the analysis and knowledge exploration of listening histories. This was complemented with brushing and highlighting techniques to uncover listening trends about artists, albums and songs. Experimental evaluation with users revealed that they were able to complete all the requested tasks with a low error rate, and that they found the solution flexible and easy to use. Additionally, users were able to infer about their main life events and listening changes, which indicates that our combination of visualization techniques is effective in conveying relevant information about the listening habits.

## Categories and Subject Descriptors

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

## General Terms

Design, Experimentation

## Keywords

Listening History, Interactive Browsing, Visualization

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AVI '12, May 21-25, 2012, Capri Island, Italy.

Copyright 2012 ACM 978-1-4503-1287-5/12/05... \$10.00.

## 1. INTRODUCTION

Nowadays, due to the proliferation and easy access of purpose specific online lifelogging services, millions of people spend their time recording facts about their lives. Exercise and running habits<sup>1</sup>, mood<sup>2</sup>, or even their music listening histories<sup>3</sup> are just a few examples. This data can be used for different tasks, such as, creating personalized recommendations, user profile characterization, or even pattern detection in their habits. Nevertheless, visualizing this data is also extremely important, because it can help users to obtain insights on intrinsic information more easily, by using a graphical representation of these facts [23].

Although these services often enrich the users' data with statistics and small graphics, their main use still is to record and to allow direct access to the information, raising the opportunity to create visualizations for this type of personal information. Regarding music listening history, some solutions and visualization techniques have already been developed, being the most relevant those developed by Byron and Wattenberg, the *StreamGraph* [5], and those created by Baur et al, *LastHistory*[4], and the *Tangle, Strings and Knots*[3]. Other community-created static visualizations, such as the *Scrobbling Timeline*<sup>4</sup>, or the *Last.fm Explorer*<sup>5</sup>, 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 design and aesthetics, and only provide static visualizations and overviews of the listening histories; poor or non-existent support for interactive browsing and filtering; and finally, scalability issues, regarding the number of songs to represent.

In this paper we describe MULHER, a novel approach for interactive exploration of music listening histories. Our solution collects data from the *Last.fm* service to create a combined visualization using a zoomable timeline and a set of synchronized-views (see Figure 1). This solution provides not only an overview of the entire listening history, but also an interactive mechanism to inspect relevant time periods or specific *data elements* (artists, albums, songs). Users can combine information from different views to unveil trends on

<sup>1</sup><http://www.nikeplus.com>

<sup>2</sup><http://www.moodstats.com>

<sup>3</sup><http://www.last.fm>

<sup>4</sup><http://playground.last.fm/demo/timeline>

<sup>5</sup>[http://alex.turnlav.net/last\\_fm\\_explorer/](http://alex.turnlav.net/last_fm_explorer/)

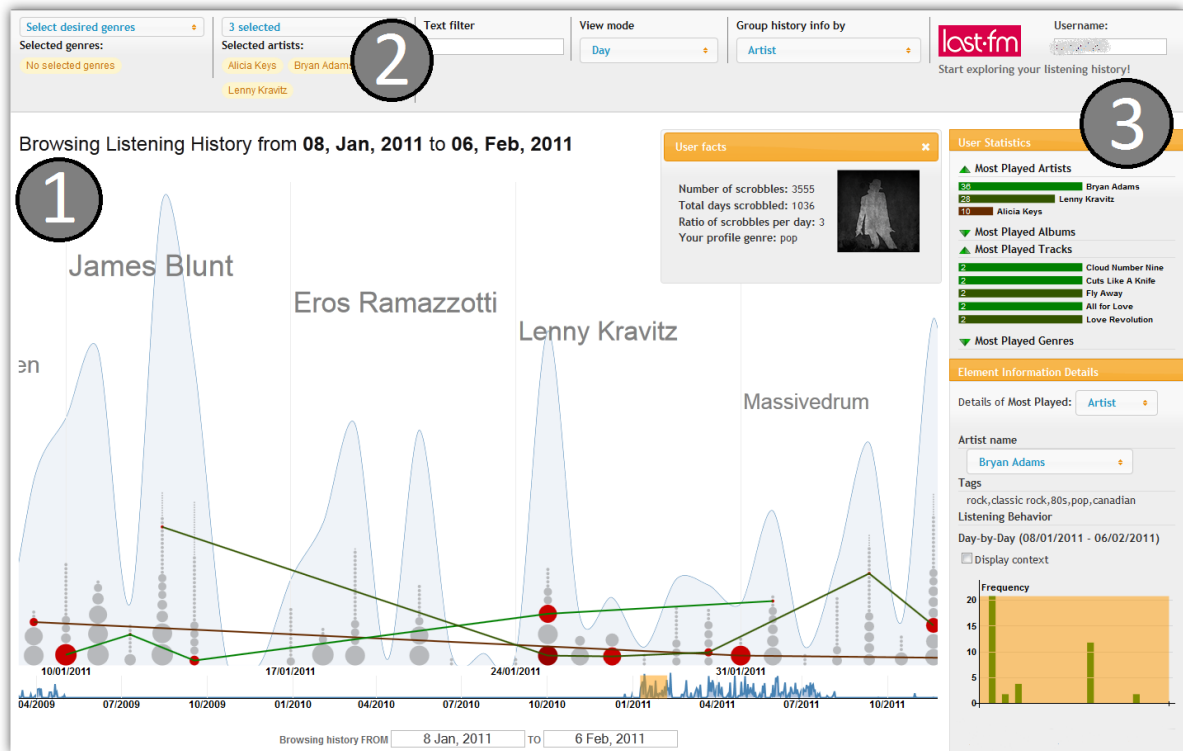


Figure 1: User interface prototype.

their listening habits, and remember facts about their lives. MULHER supports listening histories up to thousands of songs, and was built as a web application, without requiring any external software plugin.

Although the visualization techniques used in our solution are not new, our main contribution relies in being able to combine them to work well together, for the analysis and knowledge exploration of the listening history. Based on this combination we produced a tool where browsing and filtering mechanisms are flexible and effective, allowing users to explore various features at the same time. Remaining contributions include: the zoomable timeline mechanism, which avoids visual clutter (*effectiveness*) and gives users total control to browse the listening history (*flexibility*); and the new feature introduced, the "age" of the songs, that allows users to uncover new listening patterns and characterize them.

To evaluate our solution, we conducted two experiments: the first using a common listening history, and the second using personal listening logs. Overall results showed 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. Users also described our solution as very flexible and that it provided an engaging and rewarding experience. Results from the second experiment showed that users were able to recall facts about their lives, either work or personal related events, mainly by looking at the genres they were listening to, which artists/songs were played the most, or even if no music was heard at all. The interactive browsing and filtering mechanisms were considered as the

greatest advantage of our approach.

In the remainder of this document, we start by presenting the related work for our solution. Then, we describe our approach and its design rationale. 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.

## 2. RELATED WORK

To design and build our solution we reviewed work on various research fields.

First we studied information visualization, mainly temporal and timeline techniques, browsing and exploratory search interfaces, and web-based visualizations, as background for developing our tool.

Next we analyzed existent techniques and approaches to visualize and browse listening history information, identifying their contributions and limitations as the starting point for developing our solution.

### 2.1 General Information Visualization

One of the most commonly researched data types is temporal and timeline data, and various approaches to visualize it exist (see [2] and [14] for an overview). Although application scenarios often involve spatial and scientific information, there are also visualizations for non-spatial and biographical data, such as the *LifeLines* [15], a prominent

example for general-purpose creation of personal timelines in medical or legal domains (which was extended by *PatternFinder* [7] that focused on event-based patterns).

Techniques for browsing and searching in media collection have been researched too. Photo browsers often contain zoomable timelines (such as Time Quilt [10]) or focus on displaying representative photos chronologically (Calendar Browser [8]). MusicLand [11] and MuVis [6] are two approaches that support exploratory search within music collections. The Disc visualization [19] provides an overview and access to large collections of that kind. Still, personalization is not in the focus of these approaches and neither usage nor browsing histories are incorporated. MULHER integrates a zoomable timeline to interactively explore the listening history, and also to provide an overview over the user's entire listening history.

Another *chronological* type of personal data are e-mail conversations. The work by Viégas et al. partially inspired MULHER: PostHistory [20] and Themail [21] provide chronological insights into communication patterns. *Reminiscing* and making sense of the past are central activities in both approaches. Previous web-based visualization systems include Many Eyes [22], a website which visualizes user contributed data in a number of selectable visualizations; Name Voyager [25], a site that allows visualization of baby names over the last 100 years; and sense.us [9], a site that facilitated collaborative visualization, discussion, and annotation of U.S. Census data.

## 2.2 Listening History Visualization

Over the past few years, some solutions and visualizations have been developed to specifically visualize listening history information. Byron and Wattenberg, back in 2008, developed a kind of stacked graph, called *StreamGraph* [5], to visualize trends in personal music listening, where layout emphasized legibility of individual layers, arranging the layers in a distinctively organic form. Although not created to be interactive, this technique had the ability to communicate large amounts of data to the general public in an intriguing and satisfactory way. Later, some fan-created static visualizations were created, ranging from timelines displaying the number of logged songs (Scrobbling Timeline), or arcs diagrams [24] displaying how frequently listening habits change (or how mainstream specific music tastes are); to others based on stacked-graphs too, like the *Last.fm Explorer* [16] or the LastGraph<sup>6</sup>, where interactivity and support for exploration within restrictive perspectives was added.

Recently, Baur and Buts [3], using a force-directed node-and-link diagram and some variations, developed three track-based visualizations for this kind of data and proposed an automatic *playlist* generation mechanism. Although suffering from scalability and interaction problems, the authors introduced pertinent concepts. Later, Baur et al. went further by clearly exposing the issues and problems of dealing with music listening data, such as, real-world constraints, the use of contextual information and tasks users can perform using this data [17]. They developed LastHistory [4], a desktop solution that uses context information from per-

sonal calendars and photographs. However, they found that in some cases, people tended to obtain insights only based on time.

## 3. DESIGN AND RATIONALE

### 3.1 Objectives

The main objective of our solution is to facilitate exploration and visualization of listening histories in a way that can lead users to reason about their listening habits. To achieve this goal, we developed a rich-featured timeline-based visualization that allows users to interactively explore their listening history in three different perspectives: *artists*, *albums* and *tracks*. Combining this technique with a set of synchronized-views and a dynamic filtering mechanism, we expect that users will be able to answer the following questions:

- *Exploration*: What was the period when I heard more songs? Which were my most played artists, albums and songs? From which genres? How did my listening evolved overall?
- *Inference*: Do I usually listen to complete albums or prefer to hear songs from different artists? Are my favorite artists/songs the ones I heard the most over time? Am I a nostalgic person or do I prefer listening mostly to recent music? How did my tastes (genres, artists, etc.) changed over time? Why?

### 3.2 Visual Design

Listening History data is inherently based on time, as it represents the sequence of songs listened by a user over time. To reflect this structure, our solution arranges visual items on a timeline-based visualization (see Figure 1). The main visualization area (1) contains the *timeline* and the *stacked dot* visualizations for exploration. The timeline, besides providing an overall perception of all the listening history, guides the user while exploring his/her listening history. In the dotted visualization, much like what Viégas did in [20] and [21], each *stacked dots* bar encode a time interval of the selected time period, and each dot represents a *data element* (artist, album or track) listened in that interval. For example, in Figure 1, each column represents a day, and each dot an artist listened in that day.

The background of each time period is fulfilled with vertical lines that reflect the major time intervals, labels of the most played data elements, and a graphical representation of the overall listening frequency in that intervals. Although this decision was originally made for aesthetic purposes and to contextualize the intervals, experimental evaluation showed that the background also acted as a visual hint to help users exploring the information.

The size and order of the dots encode the listening frequency of the data element, and its relevance in the full listening history. The color of the dot, unlike what other solutions do, introduces a new concept to characterize the users habits: it represents the *age* of the songs, how recent or how old they are, based on how distant the release year is from the year when they were listened. The colors used to encode this feature range from strong and warm ones (e.g. red tones), representing recently release tracks (or albums), to colder colors (e.g. blue tones), representing older elements. This property

<sup>6</sup><http://lastgraph.aeracode.org/>

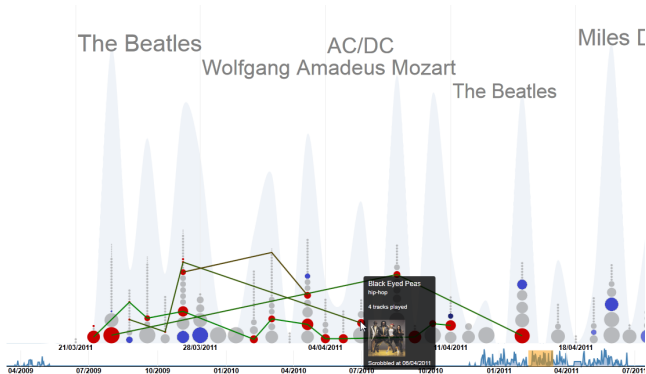


Figure 2: Multiple selection in main visualization.

allows users to get insights about their listening habits, for example, they can observe their preference for recent songs over old ones, or changes in listening to songs/albums from the same artist but from different decades, something that cannot be answered using current state-of-the-art solutions.

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* (for the remaining metadata). 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 as further described.

Statistics panel (3) displays general statistics about the most played artists, albums, tracks and genres, in the selected time period. By default, the panel displays the top five elements, but by demand, users can obtain information for less/more elements than the initial value. This data is represented by using 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, as explained in the next section.

### 3.3 Interaction

Interaction in MULHER was designed based on heuristic guidelines from information visualization: "overview first, zoom and filter, then details-on-demand" [18]. The developed components also serve as a dynamic query interface [1], that allow users to create queries dynamically and get instant feedback by *brushing and highlighting data elements*, adjusting the *time slider* or applying filters from the filtering panel. It is proven that highly interactive interfaces engage users in performing exploratory tasks [13], and therefore fit the characteristics of browsing listening habits.

Our tool starts by providing an overview of the *last month* listening history. It displays the most played songs, artists, albums and genres in this period and by default, details about the first most played elements. Starting here, users can begin exploring the history, either by changing the time period (using the slider or calendars), apply filters, or inspect

the details about the data elements.

To avoid visual clutter and enhance the exploration experience, the timeline uses an adjustable algorithm, that based on the length of the selected time period, changes the group mode to the most suitable one. For example, in a day-by-day visualization if the dots size goes below a given threshold, the group operation is changed to one higher level in the temporal hierarchy (e.g., from day to week, week to month and month to year). This feature enables users to interactively browse the listening history by time, always using the best suitable visualization over data. However, users always have control over the applied zoom and can switch between group modes whenever they need more precise details.

To allow users to get additional information and track trends on songs, albums or artists listening, *highlighting* and *brushing techniques* were used. When the user highlights a dot, additional information is displayed about the encoded data and a line appears connecting all the dots that represent the same *data element*. Connecting the dots acts as visual clue to help users tracking trends on artists, albums and songs listening. Although using this approach an increase in the line slope relates to a decrease in frequency, experimental evaluation showed that it did not mislead users understanding the information. When clicking over a dot (or several, pressing the *CTRL* key), a filter is applied by adding new criteria to the filtering mechanism (see Figure 2).


As for filtering, users can limit their search space by removing irrelevant or uninteresting items, and applying the available filters. *Genre* filter makes the dots in the visualization to *fade* if they encode data from another genre, or to maintain its color otherwise. This method gives users the ability to infer about the listened songs in that period and that genre: "Were they recent or old?"; or follow the listening trends of that genre over time. Artists (and text filter), make the dots to be connected in the main visualization, as described in the previous paragraph.

In the details panel, users can get details on metadata information, but also on day-by-day and hourly playing information (see figure 3). The day-by-day visualization consists in a *bar chart* where each bar represents a day in the selected time period (see figure 3a), and the hourly visualization is a *dot chart*, where each dot represents an element listened at a specific time [day, hour-minutes]. These visualizations can display the overall listening history for a selected *data element* (artist, album, song), while highlighting the selected time period (see figure 3b). The analysis of these two visualizations can help users to obtain insights about how they listened to that particular *element*, answering questions like "Which days did I listen to Colbie's music (and how much) and at what time? Do I hear it a lot? Always at morning?".

## 4. EXPERIMENTAL EVALUATION

In order to evaluate our approach we conducted two experiments with users.

First we intended to evaluate the users' ability to complete all the proposed tasks, measure the error rate, and their satisfaction and experience, while browsing a common music listening history. The second experiment was designed to



(a) Day-by-day with no context (b) Hourly with surrounding context

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

evaluate if users could use our solution to explore their own listening histories, and get insights on their listening habits and behaviors.

For both experiments, the procedure and tasks used were the same, but the listening logs were different. In the first test we used the same log for all users and in the second we used the users' logs.

## 4.1 Users

For the first experiment, we conducted tests with 10 users (2 females and 8 males), with 7 of them aged between 20 and 30 years, and the remaining between 30 and 50 years. 7 out of 10 users 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). We asked about the places they usually listen to music, 9 said they listen to music at home, 8 at work and while driving. All the users answered that they listen to music in computer, 8 of them listen to radio (streaming or regular), and 7 use smartphones to listen to music when moving around.

In the second evaluation, tests were conducted with 4 users, all male and aged between 20 and 40 years. Two of them were computer science assistant teachers, one a senior software engineer, and the last one an undergraduate PhD student. The 4 users listened to music every day and had a *Last.fm* account. Listening histories ranged from 3.000 to 30.000 songs.

## 4.2 Apparatus

MULHER prototype was implemented as a regular web application, using HTML, CSS and Javascript. The visualizations were created using Protovis<sup>7</sup>, a graphical library for creating custom visualizations. A Java backend was developed in order to store and pre-process the information collected and to improve the tool performance. JSON was used as the standard format to represent all data, as it is natively used in Protovis (and javascript).

The listening log used in the first experiment was collected from the *Last.fm* account of one of the authors, to setup a

<sup>7</sup><http://mbostock.github.com/protovis/>

baseline for all the test sessions. It contained approximately 5.000 records from November 2009 to July 2011. This user regularly scrobbled his listening history, mainly for the last year.

## 4.3 Procedure

Each test took around 45 minutes, and was performed using users' personal computers. The only requirement about the setup was that an internet connection was available.

First, we started with a brief description about the experiment, thanking the person and explaining our objectives and what we were going to do. Next we introduced our application, describing its main features and how to use them. At this point we gave a practice time for users to get familiar with the solution. A video was used to present the interaction and main features of the solution. After this introduction, users performed a set of tasks using the prototype. Tasks covered different objectives as explained below. When authorized, users interaction with the application was recorded, using a screen capture software for later post-analysis. Finally, users fulfilled a satisfaction questionnaire and an informal interview was conducted to get users' general impressions about the experience.

## 4.4 Tasks and Objectives

Based on the analysis of the tasks presented in related work, mainly in [4], and the tasks described by Sellen and Whittaker in their recent overview of user tasks in lifelogging applications[17], we designed a set of tasks to evaluate our solution, focusing on *reminiscing*, *retrieving*, and *remembering* intentions. The tasks are summarized in Table 1.

The proposed tasks were divided in two groups accordingly to our main objectives. In the first group, we tried to evaluate if users could use our approach to effectively explore and browse the listening history. Tasks ranged from direct inspection of the solution's interface (e.g., getting information about most played artists, or the period with more *intense* music listening activity), to others more complex, where users had to combine browsing and filtering (with genre and artist information) to obtain the correct answers for the proposed questions. 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 a song listening evolution (e.g., if it increased/decreased or was constant, etc.), or trying to understand some aspects of the user's profile, such as, "*Does he/she listens to complete albums, or he/she prefers to hear different songs?*".

## 4.5 Results

### 4.5.1 Completeness and Error Rate

In short, all users from both tests were able to complete the required tasks in the given time period. Users did not find any relevant difficulty in executing the tasks.

Overall, 13 of the 14 users found it easy to execute the required tasks (see Table 1), and only user 1.1, in task 5, found it to be somewhat difficult (although he classified it as *easy*). He reported that the brushing line slope confused him mo-

Task	Description	Very Difficult	Difficult	Normal	Easy	Very easy
1	Indicate the date of the first song listened by the user.	-	-	-	60%	40%
2	Indicate the date of the last song listened by the user.	-	-	-	60%	40%
3	Indicate the month where more songs were listened.	-	-	-	60%	40%
4	Find the most played song of the artist "Bryan Adams" over the last three months.	-	-	-	10%	90%
5	Describe the trend on the previous identified song.	-	-	-	10%	90%
6	Indicate the most played songs, albums, artists and genres.	-	-	-	10%	90%
7	For each of the identified artists, indicate if more individual songs or complete albums were heard.	-	-	-	20%	80%
8	Describe and try to justify the listening changes that occurred over the last three months (if any exists).	-	-	10%	80%	10%
9	Describe the listening habits on the selected period. (in terms of the age of songs, time, artists played, etc.)	-	-	10%	70%	20%

Table 1: Tasks description and easiness classification results.

mentously. But after this incident, he continued the tasks execution and finished all of them without any other issues.

Task success rate was very high for all the tasks, with an average value of 97%. A task was considered successfully completed if the user finished it and gave a correct answer (when required). Indeed, only on task 3, where users had to obtain an overview of the listening history, and infer about the most *intense* listening period, we got poor results (with an error rate of 40%). This suggests that improvements should be made to provide more accurate information when answering overview issues, perhaps by completing visual representations with numeric ones.

#### 4.5.2 User Satisfaction

After tasks execution, users fulfilled a satisfaction survey. Most questions were rated using a *5-point Likert scale*, ranging from *Very Difficult* to *Very Easy*, except for the ones using semantic scales and the open ended questions.

Satisfaction survey revealed some interesting results. 13 of the 14 users were able to use the application and understand the information presented. 4 of them felt unsure while executing task 5 (describing a song's listening history), mostly because they were afraid of giving a wrong answer, as this was a subjective and interpretative question.

As we intended to assess the overall user experience of using our solution, we asked users to rate it using a different set of semantic scales. Results indicate that 12 of them found the experience *engaging* and *rewarding*. Users expressed their joy of learning something about the other user's listening history and their own, as reported in the second experiment.

All users said that it was important to understand the listening history, but only *three* have previously tried it. These ones 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 very easy to obtain it using our solution. When inquired about using this information to improve recommendation systems, for example, to suggest songs or help automatic *playlist* generators, they all agreed that this is indeed crucial. However, they don't

understand why most available personal players do not make use of it, mostly because this data can be implicitly recorded by applications and through user interaction with them.

In general, users enjoyed using our solution and they intended to use it in their daily routine, to understand more about themselves and even to share their habits with others. One user even said: "*This tool is very nice. I would like to use it to share my listening history. I wonder which patterns are shared among my friends and those I can see using the tool!*".

#### 4.5.3 Inferring and detecting listening patterns

*Reminiscing* and *remembering* are central tasks to lifelogging applications, and some of the proposed tasks intended to capture how this could be accomplished using our solution. The tasks 5, 7, 8 and 9 required users to perform some interactive exploration over different views and perspectives to answer them.

In the first experiment, although users were browsing another person's history, they all were able to understand four aspects of the user habits: recently, he played much more songs from other artists, than songs from his favorite ones; he prefers to listen to complete albums almost every working day (in the morning and afternoon), but listens to diverse songs during the night; there was huge seasonal changes in the genres the user listened to (changing from *Pop/Rock* to *Reggaeton* in the summer, and then passing through *Jazz* and ending again in *Pop/Rock*); he prefers recent music (perhaps by seeking the tops), but some days, "*looks like he gets nostalgic, because he only hears old songs.*", as one subject said. This could not have been realized by using any of the solutions describe in related work.

However, in first experiment, we could not evaluate if users could remember past events, since it was performed using another person's listening history. Giving users the opportunity to browse their own listening histories (as in the second experiment), made it possible to evaluate if they could find personally relevant listening patterns and relate them with their daily life routine.



All users were able to discover something new about their listening habits, understand intrinsic information contained in the log, and that they were not aware. Most of this insights were based on the absence of music listening, trends on artists, songs and genres listening. For example, user 1.1 was able to identify different types of stuff he did at his work, based on the timeline and the different overviews our tool provide; he said: *"Here I was working on a scientific paper, because I was listening only classical music, and I like to hear that kind of music when I'm writing."*, and then he added, *"But then I skipped listening to music, because I had some project discussions, and no time to listen to music"*. User 1.2 was able to confirm that he prefers to listen to more old music, only by looking to the overall color represented in the *dotted visualization*, and he navigated to a specific period and asked the interviewer: *"Well, do you know why this part of the visualization contains mostly recent music, even though I just prefer to listen old music?"*, and then he answered that the reason was that one of his favorite artists just released a new album, after years of absence in musical scene. None of the solutions described in related work section could help him understand this fact.

Time patterns during the period of the day was also detected by two users. User 2.4 reported that he was not consciously aware that he was listening to much music until late in night, he said: *"I did not realized that I was listening too much music in late night"*, and then continued *"but now that I think of this, I usually listen to more rhythmic music at that time to stay awake a little longer, mostly when I'm working"*. User 2.3 was even more thoroughly descriptive and said *"Looks like that through a regular day I keep changing the genre of music I listen to. Start with something stronger in the morning and then end the day with more relaxing songs!"*. User 2.3 was also surprised when exploring his information and found a week where he listened to a *Jazz* artist, which he does not even remembered he had. However, by looking at the time he listened to it and the albums and songs from that artist, he remembered that on that day he went to a concert of that artist at night. This was purely based on time and intrinsic information contained in the listening history, not requiring any external/contextual information like in [4].

Regarding trends on artists, albums, songs or genres listening, user 2.1 and 2.4 reported some facts. Based on the visual hints provided by the *background labels*, and while inspecting a period that caught the user's attention, user 2.1 commented that the *"lines connecting the dots were very useful to understand how the listening of a specific artist/song evolved over the time"*; user 2.4 added: *"I've found songs that I only listen in the weekend, because I'm home by that time; but it is interesting to see that, during the whole week, I can observe changes on the listening frequencies of different artists, because, even though my tastes remain the same, they change over the week and the places where I am."*

## 5. DISCUSSION

Our solution was designed based on the hypothesis that providing an interactive browsing and exploration visualization tool would allow users to browse their listening history and make sense of the past, identifying some patterns hidden in the listening history and reasoning about it. As presented in

the previous sections, results from the performed evaluation supported this hypothesis.

The *timeline-based* mechanism and the adjust algorithm used to narrow the *stacked dot* visualization, proved to be a major asset of our solution, not only because they stand out as our main browsing and filtering techniques, but also because their *effectiveness* and *flexibility* were validated by the results obtained from the experimental evaluation.

Although designed for aesthetic reasons and provide context information of the listening history, different overviews and background information, ended being an important aspect of our solution. It not only contains *context information* of most played artists or songs in a selected period, but also helped users selecting a start point to explore the information (acting as visual clues/hints to choose a start point to begin or continue exploring).

Results from experimental evaluation, showed that the new feature described in this paper, the "age" of the songs, can effectively transmit to users information about their listening habits. Different listening profiles can be detected by direct color inspection, and indeed useful for users comparison or similarity inference.

Knowledge analysis and inference could only be performed by combining all the visualizations and obtain information from different perspectives of the same data. As an example, combining temporal filtering with *keywords*, with the day-by-day and hourly visualizations and the "age" of songs, some users could observe personally relevant patterns on their listening habits and behaviors, and answer questions like the ones proposed in the objectives. Different patterns were possible to be identified, being the most relevant ones: regular changes in the listened genres; trends on artists/songs listening frequency; different listening periods during a day to listen to different artists/songs; and finally, understand users preferences for recent or old music. *Reminiscing* and *remembering* was even possible for some users, that described some life events and their daily routine mainly based on time, absence of music listening and the context information provided by the visualizations.

In summary, the combination of different visualizations and the features used to browse and obtain information, stand out as our major contribution and the main asset of this work.

## 6. CONCLUSIONS AND FUTURE WORK

In this paper we described a novel solution for browsing and exploring music listening histories. The solution combined a timeline-based visualization with a set of synchronized-views to allow users to perform an interactive exploration of their music listening histories. We also introduced a new feature, the "age" of songs, to characterize and encode the color of graphical elements in the *stacked dot* visualization. Results from user evaluation showed that the developed solution was easy to use, with users completing all the tasks with a low error rate, and founding the experience engaging and rewarding. Experiments with users' personal listening histories showed that listening patterns and life events can be remembered and contextualized by using our approach,

and only based on the listening history itself. Changes in the listened genres or the absence of listening and their correlation with users' daily routine, are two examples of patterns detected.

Regarding future work, we are planning to explore new and different perspectives over data to perform browsing. We intend to use data mining algorithms to automatically discover listening patterns, and then graphically represent this information, in a way that can not only show new insights about users' listening habits, but simultaneously guide them in exploration tasks. As an example, this information can provide insights on a set of user profiles: users that always seek the tops; others that enjoy mostly female voices or that listen to classic music in the morning but rhythmic music in the afternoon; etc. Finally, these new perspectives can then be used for music recommendations and automatic *playlist* generation within different contexts and activities, like the ones described in[12].

## 7. ACKNOWLEDGEMENTS

This work was supported by FCT (INESC-ID multiannual funding) through the PIDDAC Program funds and through the Crush project, PTDC/EIA-EIA/108077/2008.

## 8. REFERENCES

- [1] C. Ahlberg, C. Williamson, and B. Shneiderman. Dynamic queries for information exploration: An implementation and evaluation. ACM Press, 1992.
- [2] W. Aigner, S. Miksch, W. Müller, H. Schumann, and C. Tominski. Visualizing time-oriented data - a systematic view. *Comput. Graph.*, June 2007.
- [3] D. Baur and A. Butz. Pulling strings from a tangle: visualizing a personal music listening history. IUI, 2009.
- [4] D. Baur, F. Seiffert, M. Sedlmair, and S. Boring. The streams of our lives: Visualizing listening histories in context. *IEEE TVCG*, 2010.
- [5] L. Byron and M. Wattenberg. Stacked graphs geometry & aesthetics. *IEEE TVCG*, 2008.
- [6] R. Dias and M. J. Fonseca. Muvis: an application for interactive exploration of large music collections. In *Proc. of ACM Multimedia*, 2010.
- [7] J. Fails, A. Karlson, L. Shahamat, and B. Shneiderman. A visual interface for multivariate temporal data: Finding patterns of events across multiple histories. In *IEEE Symposium On Visual Analytics Science And Technology*, 2006.
- [8] A. Graham, H. Garcia-molina, A. Paepcke, and T. Winograd. Time as essence for photo browsing through personal digital libraries. In *Proc. of the second ACM/IEEE-CS Joint Conference on Digital Libraries*, 2002.
- [9] J. Heer, F. B. Viégas, and M. Wattenberg. Voyagers and voyeurs: Supporting asynchronous collaborative visualization. *Commun. ACM*, 2009.
- [10] D. F. Huynh, S. M. Drucker, P. Baudisch, and C. Wong. Time quilt: scaling up zoomable photo browsers for large, unstructured photo collections. In *CHI Extended Abstracts on Human factors in computing systems (CHI EA '05)*, 2005.
- [11] H. Lam and T. Munzner. Increasing the utility of quantitative empirical studies for meta-analysis. In *Proc. of the conference on Beyond time and errors: novel evaluation methods for Information Visualization (BELIV '08)*, 2008.
- [12] D. J. Levitin and J. McGill. Life soundtracks: The uses of music in everyday life'. *unpublished Dept of Psychology McGill University*, 2007.
- [13] G. Marchionini. Exploratory search: from finding to understanding. *Commun. ACM*, 49:41–46, April 2006.
- [14] W. Muller and H. Schumann. Visualization methods for time-dependent data - an overview. In *Simulation Conference*, 2003.
- [15] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman. Lifelines: visualizing personal histories. In *Proc. of the SIGCHI conference on Human factors in computing systems: common ground (CHI '96)*, 1996.
- [16] M. A. Pretzlav. Last.fm explorer: An interactive visualization for hierarchical time-series data. *University of California, Berkeley*, 2008.
- [17] A. J. Sellen and S. Whittaker. Beyond total capture: a constructive critique of lifelogging. *Commun. ACM*, 2010.
- [18] B. Shneiderman. The eyes have it: a task by data type taxonomy for information visualizations. In *Proc. of IEEE Symposium on Visual Languages*, 1996.
- [19] M. Torrens and J. lluis Arcos. Visualizing and exploring personal music libraries. In *In Proc. of the 5th International Conference on Music Information Retrieval (ISMIR '04)*, 2004.
- [20] F. Viegas, D. Boyd, D. Nguyen, J. Potter, and J. Donath. Digital artifacts for remembering and storytelling: posthistory and social network fragments. In *Annual Hawaii International Conference on System Sciences*, 2004.
- [21] F. B. Viégas, S. Golder, and J. Donath. Visualizing email content: portraying relationships from conversational histories. CHI, 2006.
- [22] F. B. Viegas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon. Manyeyes: a site for visualization at internet scale. *IEEE TVCG*, 2007.
- [23] F. B. Viégas. Newsgroup crowds and authorlines: Visualizing the activity of individuals in conversational cyberspaces. In *Proc. of the 37th Hawaii International Conference on System Sciences*, 2004.
- [24] M. Wattenberg. Arc diagrams: Visualizing structure in strings. In *Proc. of the IEEE Symposium on Information Visualization (InfoVis'02)*, 2002.
- [25] M. Wattenberg. Baby names, visualization, and social data analysis. In *Proc. of the IEEE Symposium on Information Visualization (INFOVIS '05)*, 2005.