

Gaining Musical Insights: Visualizing Multiple Listening Histories

Ya-Xi Chen
Media Informatics
University of Munich
yaxi.chen@ifi.lmu.de

Dominikus Baur
Media Informatics
University of Munich
dominikus.baur@ifi.lmu.de

Andreas Butz
Media Informatics
University of Munich
andreas.butz@ifi.lmu.de

ABSTRACT

Listening histories are rich sources of implicit information. Combining listening histories of multiple users allows comparisons in musical taste and discovery of new music, though few existing work specifically addresses this issue. In this paper we present two interactive visualizations which give users a deeper insight into consent and dissent in their listening behaviors, and help them to compare their musical tastes. HisFlocks shows overlaps in genre and artist in certain time periods and LoomFM illustrates sequential listening patterns. Our first feedback on these systems was very promising and we plan to extend our concepts to broader scenarios with a greater number of listening histories.

Author Keywords

Listening history, visualization, multiple users, Last.fm, time-based visualization.

ACM Classification Keywords

H.5.2. User Interfaces: Graphical User Interfaces (GUI).

INTRODUCTION

"What kind of music do you like?" is a common question while getting to know someone. Learning about musical taste is appealing and there are various effects of social bonding surrounding music [2, 4]. Still, describing one's musical taste in a short and precise way is hard, so mostly the above question is answered vaguely, for example by "I like all kinds of music". Online services like Last.fm can help here: The accompanying Audioscrobbler plugin automatically logs all music that the user listened to and publishes it in his/her profile page, thus making music consumption social without additional effort. This personal profile with lists of tracks recently listened to and top art-

ists/tracks in different time periods provides an overview of one's taste in music. Last.fm's large collection of data also allows automation of social music tasks: Using collaborative filtering, recommendations of similar music can be generated from the user's profile and users with similar taste (so-called "neighbors") can be determined. Still, Last.fm's web-bound services are mostly textual and limited (for neighbors, overlaps in artists is all that is given as information on similarity). To grant users access to this vast data-set, graphical visualizations are needed to let them understand personal listening behavior, re-discover forgotten music and understand others' taste in music.

In this paper, we explore different interactive visualization concepts for multiple users' listening histories and propose two prototypes which facilitate the comparison of multiple users' musical tastes.

There are some existing visualization tools regarding time-based data and specifically listening history. Time-based data is mostly visualized as a temporal flow which indicates hidden trends over time. ThemeRiver [6] introduces a river metaphor to represent topical changes within a large collection of documents. Similarly, most of the existing visualizations for personal listening histories use the flow metaphor. Extra Stats¹, LastGraph² and Last.fm Spiral³ all extract artist or genre trends from consumed songs and show their impact as waves along a time-line. In our former work [1], we deviate from this trend by visualizing listening histories either as chains of songs based on self-contained listening sessions or based on listening proximity. Last.fm Explorer [7] is one of the few works that allows showing two users' listening histories as two respective graphs (see Figure 1) which can be displayed on 3 different levels based on Last.fm tags, artists or tracks.

In order to gain more insight into the effective use of metadata contributed by the users, we discussed with 13 Last.fm users about the issue of musical taste and listening history. During the interview, the participants were shown the

Workshop on Visual Interfaces to the Social and Semantic Web (VISSW2010), IUI2010, Feb 7, 2010, Hong Kong, China. Copyright is held by the author/owner(s).

¹ <http://build.last.fm/item/34>

² <http://lastgraph3.aeracode.org/>

³ <http://build.last.fm/item/377>

Last.fm Explorer visualization result of their own listening history and one of another participant. The concept of visualization of multiple users' listening histories received positive feedback. It was appreciated that the history can be displayed in three different levels. Participants appreciated that they could control the visualization with interactions such as filtering, stacking and brushing. But their comments revealed that Last.fm Explorer's focus on a single user's listening history leads to some intrinsic issues in need of improvement for visualizing multiple histories: As shown in Figure 1, only a recent part of history can be shown in the limited space, so the two timelines might be misaligned especially when one user did not listen to music at the beginning or end of the dataset. Although information can be displayed on three different levels, the lack of zooming might cause users to lose the current context. It is not possible, for example, to learn an artist's genre by switching from artist to genre view, because the size and placement of items change. The main issue of Last.fm Explorer claimed by the participants, however, is that it is difficult to actually compare user's taste in music: The stacked graphs make the impact of certain artists or tracks hard to read and finding correlations in two graphs can only be performed by a lengthy visual search process.




Figure 1. Last.fm Explorer: visualizations for two users.
http://alex.turnlav.net/last_fm_explorer/

Although existing visualization tools receive positive feedback, most of them target single users and present static, non-interactive images. There is a lack of tools for browsing and comparing multiple users' listening histories directly instead of tediously clicking through web pages.

VISUALIZATION OF MULTIPLE USERS' LISTENING HISTORIES

Based on the analysis of existing tools and results of the interview, we gathered the following design principles for visualizing multiple users' listening histories within an interactive interface:

Time period consistency: Since a listening history is time-based, the visualization should be built with a global timeline. Each listening history should be aligned to facilitate comparisons within the same time period.

Good overview: The visualization should offer a good overview especially for long listening histories, so the user can get a general impression at first glance without continuously switching back and forth.

Managing complexity: Music-related information is complex and has metadata on various levels. The most prominent ones are genre, artist and track. Combined with information about the consumption behavior, such as listening histories, the complexity steadily increases. Therefore it is essential to support the user's task specifically comparison by focusing on key data elements and controls.

Easy to compare: Multiple listening histories should be represented in one interface to facilitate convenient comparison between different users.

Based on these design principles, we developed two different visualizations, Last.fm HisFlocks and LoomFM, aiming at different aspects: HisFlocks gives users an overview of relations in genre and artist in certain time periods, while LoomFM allows a comparison of consumption patterns over time. We will present these two applications here.

Last.fm HisFlocks

Since the genre and artist level provide a better overview than the track level, we chose genre and artist as the main categories and visualized this general information in our first interface Last.fm HisFlocks (see Figure 2).

In HisFlocks, the entire history is represented as a series of time frames with all the artists listened to in each time period being displayed on a genre map, where similar artists are grouped into clusters, labeled with genre names. By default, the time interval between each frame is one week, but this can be adjusted freely. Within a frame, each artist's name is color-coded (1 color per compared user) and its size represents the play count. Shared artists are further grouped into highlighted sub-clusters. The user can pan and zoom to navigate within the map. In different frames, the same artist always appears at the same position. As Figure 2 shows, an overview for all the artists listened to is provided in the first frame and by dragging the slider, the user can browse all frames sequentially in which the artists being listened to in this time period are highlighted and others fade out. The user can easily see the change of an artist, for example, by appearance, disappearance and growth of size, while she or he adjusts the time slider, and thus have an understanding of how the musical taste changes over time. In Figure 1, the two users share a lot in *metal*, which is scattered in some sub-genres such as *heavy metal*, *gothic metal* and *melodic death metal*, and thus cannot be derived at first glance. Therefore, some aggregation of these tags could improve the understandability of musical similarity on the genre level. In HisFlocks, all the artists in one frame are

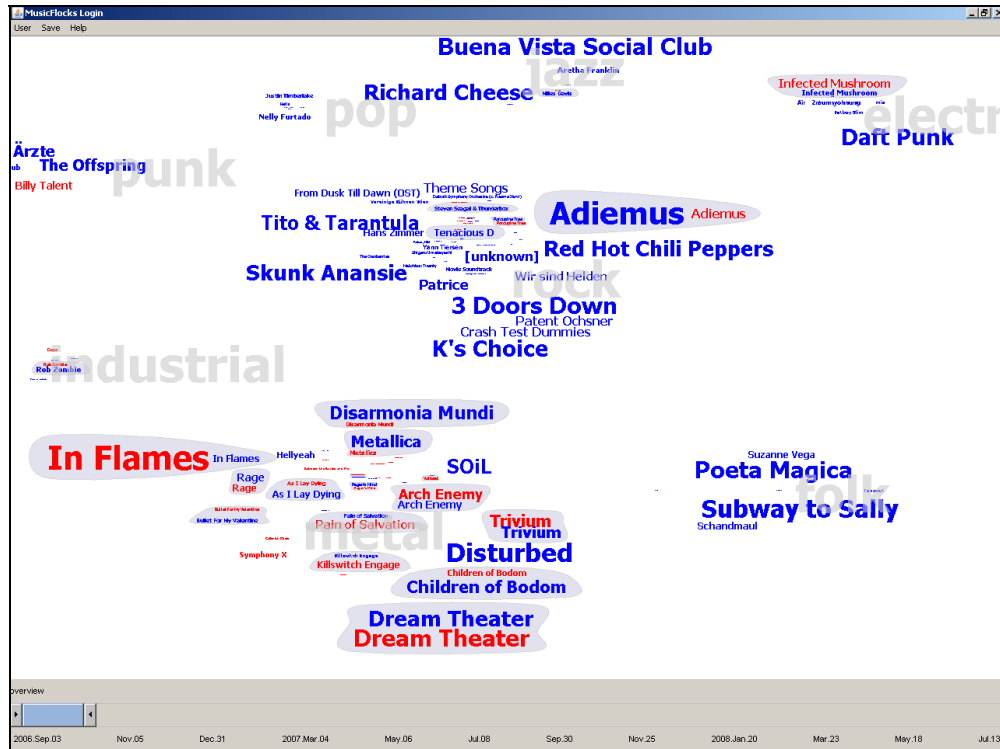


Figure 2. Last.fm HisFlocks.

grouped into genre-clusters based on their most representative genre-tags. This organization is based on a semantic analysis which determines the structure and position for each genre and artist. We applied text analysis to create a semantic clustering of similar tags while excluding redundancy: After removal of different separators, such as “_” and “&”, the Porter algorithm [8] is applied to detect the stem for each tag. Tags with the same stemmed words will be clustered in the same group: The aforementioned metal-related ones are grouped into one *metal* cluster, thus providing an abstract overview for general users. The locations of each genre and artist are determined by the semantic similarity. This semantic similarity (*Sim*) between tag A and B equals the ratio between the number of resources in which tags co-occur and the number of resources in which any of the two tags appears (see Equation 1).

$$Sim(A, B) = \frac{|A \cap B|}{|A \cup B|} \quad (1)$$

With the semantic analysis all the tags will be well organized: the initial location of each tag is assigned by means of a 2D projection based on a multidimensional scaling of co-occurrences (see [3] for more detail).

The overview shown in Figure 2 indicates that both users listen to *rock* and *metal* a lot, and they do share many artists in these two genres. Other genres they also listen to are *punk*, *industrial* and *electro*. Only the user with the blue color listens to *folk*, *pop* and *jazz*. Although they have many artists in common, they listen to few of them at the exact

same time period. For example, in one time period as Figure 3 shows, two users both listened to 3 common artists (*Killswitch Engage*, *Children of Bodom*, *Dream Theater*), but only one (*Dream Theater*) was exactly listened to by both users in this time period. Besides, only the “red” user listened to “Symphony X”.




Figure 3. Detailed view in one frame.

LoomFM

HisFlocks focuses on displaying artists on an abstract genre map while hiding the “sequential”, track-based aspects of listening behavior: How many songs are consumed in one session? How often are certain songs listened to? Are there repeating patterns within this list (such as albums or pre-defined playlists) and large gaps (due to vacations etc.)?

To address these issues we designed the complementary visualization LoomFM (see Figures 4 and 5). A horizontal timeline displays the temporal scope of two listening histories. Based on that, songs of one user are displayed above and of the other below this timeline and symbolized as small circles (hovering with the mouse above an icon gives more detailed information about what song it is and




Figure 4. LoomFM global view. Repetitions and main musical trends become visible (the white box shows the borders of Figure 5).

when and by whom it was listened to). The distance of a song icon to the middle timeline represents the overlap with the other history: similar songs will be shown with their titles and displayed very close to the center, while less similar songs that are only represented by their artists or even their genres are further away (with a completely unrelated song at the farthest). The up and down of the song circles thus represents the general similarity between the two users.




Figure 5. LoomFM detailed view. a) Timeline, b) User 1 history (red), c) User 2 history (purple), d) Artist names.

Instances of songs that appear more than once are connected by yellow arcs within one history as well as between histories. Similar to Arc Diagrams [10], several aspects can be visually derived from these arcs: For one, the higher an arc is, the further two instances of a song are apart, which means that in general the height of the arcs represents the consumption intervals of the same song (each instance is connected to all others). No arcs or very low ones show that a user has a fluctuating listening behavior with a low tendency to listen to the same songs in the long term (see lower user in Figure 4). Additionally, adjacent bundles of lines show repeated sequences of songs (for example, pre-

defined playlists or albums). In Figure 4 the release of a new album by a shared artist leads to repeated listening by both users (two yellow arc bundles to the left of *indie rock*).

Artist name and genre of each song are displayed beyond the user timelines. The higher the number of consecutive songs sharing the same genre or artist, the larger the label, which makes general trends appear at the maximum zoom level (see Figure 4: *indie* and *electronic* as genres for user 2 and *indie rock* as a common genre).

The user can again navigate the displayed song plane by panning and zooming. Zooming out gives an overview of the complete listening histories with the most influential genres and artists still visible.

SUMMARY AND FUTURE WORK

Music listening histories and their accompanying metadata are complex information that can bring rich insights into one's own and other's behavior. We explore the possibilities to facilitate browsing and comparing multiple users' listening histories, and present two interactive visualization approaches: To illustrate how the musical taste changes over time, both of the visualizations are time-based. They emphasize certain aspects of the music-related data in order to reduce the complexity and support visual comparison by merging the users' histories into one representative graph. Last.fm HisFlocks focuses on time periods and aggregates musical taste based on genres and artists within these periods. LoomFM is more timeline-centric and highlights sequential aspects of histories, such as repetitions and the differences in taste over time. These visualizations provide users with means to get a deeper understanding for neighboring listeners and also discover new songs.

Given that both systems are implemented as first prototypes, some issues need to be improved in upcoming versions. In HisFlocks, although the same artist always appears at the same position, the overall similarity-based placing of clusters and that of artists in each cluster is rather arbitrary. Users cannot anticipate where exactly new artists will appear. A more comprehensible placement would be appealing. LoomFM suffers from a low performance when displaying

more than a thousand songs, an issue that comes up with realistically sized collections.

We discussed our prototypes with Last.fm users and received rather positive feedback and valuable suggestions. We plan to add some extra functionalities, such as filters. In the current state only two histories are compared and we plan to examine our concepts with a far greater number and size of listening histories. We hope to gain additional ideas regarding understandability, data reliability and the human perception of musical taste. In our future work, we will extend our concepts to broader scenarios with consumption record, for example with books or photos.

ACKNOWLEDGMENTS

This research was funded by the Chinese Scholarship Council (CSC) and the German state of Bavaria. We would like to thank the participants of our study.

REFERENCES

1. Baur D., Butz A. Pulling Strings from a Tangle: Visualizing a Personal Music Listening History. *Proc. of IUI 2009*, 439-444.
2. Bentley F., Metcalf C., Harboe G. Personal vs. commercial content: the similarities between consumer use of photos and music. *Proc. of CHI 2006*, 667-676.
3. Chen Y-X, Santamaria R., Butz, A., Theron R. Tag-Clusters: Semantic Aggregation of Collaborative Tags beyond TagClouds. *Proc. of SG 2009*, 56-67.
4. Cunningham S., Reeves N., Britland M. An ethnographic study of music information seeking: implications for the design of a music digital library. *Proc. of JCDL 2003*, 5-17.
5. Dubinko M., Kumar R., Magnani J., Novak J. Raghavan P., Tomkins A. Visualizing tags over time. *ACM Transactions on the Web*, 2007, Vol11, No. 7.
6. Havre, S., Hetzler, B., Nowell, L. ThemeRiver: Visualizing Theme Changes over Time. *Proc. of InfoVis 2000*, 115-123.
7. Maxwell A. Pretzlav. Last.fm Explorer: An Interactive Visualization of Hierarchical Time-Series Data. <http://vis.berkeley.edu/courses/cs294-10-fa08/wiki/images/9/9f/PretzlavFinalPaper.pdf>.
8. Porter M. F. An algorithm for suffix stripping. *Program*, vol. 14, no. 3, pp 130-137, 1980.
9. Viégas, F. B., Golder, S., Donath, J. Visualizing Email Content: Portraying Relationships from Conversational Histories. *Proc. of CHI 2006*, 979-988.
10. Wattenberg, M. Arc Diagrams: Visualizing structure in strings. *Proc. of InfoVis 2002*, 110-116.