

---

# The Songs of Our Past: Working with Listening Histories

**Dominikus Baur**

University of Munich (LMU)  
Amalienstrasse 17  
80333 Munich, Germany  
dominikus.baur@ifi.lmu.de

**Abstract**

Music listening histories are portraits of a person's taste in music. In my research I am exploring this type of data and how user interfaces can be enhanced with it. In this Doctoral Consortium paper I describe my approach towards this goal: Statistical analysis and casual information visualizations can help in finding relevant patterns and aspects in listening histories. Making them available to regular users and asking what they learnt about themselves gives us the chance to find out more about their listening on the minute level of songs. Contextual information such as photos or calendar entries can help trigger memories. In this paper I describe the motivation and goals of my research and my current status. In the end, both the HCI community and end users can benefit from more convenient and sophisticated interfaces for this data.

**Keywords**

Music, Listening Histories, Lifelogging, Information Visualization, Psychology

**ACM Classification Keywords**

H5.5. Sound and Music Computing: Miscellaneous.

**General Terms**

Design, Experimentation, Human Factors

---

Copyright is held by the author/owner(s).  
*CHI 2011*, May 7–12, 2011, Vancouver, BC, Canada.  
ACM 978-1-4503-0268-5/11/05.

## **Introduction**

One of the many benefits that the transition from analog to digital brings is the ease of logging one's music consumption. Services like Last.fm<sup>1</sup> store detailed lists of all songs that a person heard automatically in the background without requiring any user interaction.

Work on the user benefits of such lifelogging data, however, is scarce [7]. For the creators of listening histories, ways to make use of the data exist only implicitly as suitable recommendations. Explicit ways to reminisce based on this data, make sense of it or use it while creating playlists or organizing music collections are not available. My goal is to provide users with tools that let them use their data for their own advantages.

## **Background and related work**

Several disciplines are relevant to my research: Music psychology and sociology are intimately concerned with the ways people listen to music (e.g., [6]) but on a rather abstract level: Manual data collection restricts the results to day-to-day timeframes instead of moment-to-moment as automatic logging does. Information visualization and specifically Casual InfoVis [5] describe the techniques and approaches to allow non-expert users to learn about and make sense of the large data sets that listening histories (and associated contextual data) represent. Research in recommender systems takes listening histories as input to their algorithms and is therefore interested in understanding users' behaviors (e.g., [4]), but favors automation before interactivity. Finally, the Music Information Retrieval community (see proceedings of the annual

ISMIR conferences) provides algorithms for content-based music analysis but also interfaces for listening to and organizing music. Integrating listening histories in these interfaces happens, if at all, only implicitly.

## **Statement of thesis or problem**

The main goal of my thesis is to understand what benefits can be provided for users from their listening histories beyond music recommendation. These benefits might lie in the area of learning more about themselves, reflecting about their behavior or reminiscing about their past. More tangible benefits could be the improvement of music-related applications for organizing one's music collection or creating playlists. My hypothesis is that listening histories can help in this regard.

## **Research goals and methods**

The problem of integrating listening histories into applications implies a sequential research process: My first research question is what information can be drawn from listening histories and how to get it. The second research question more specifically asks about the ways in which this information can be beneficially integrated into user interfaces.

To answer the first question I follow a threefold approach:

(1a) First, there is an extensive body of work from a variety of fields regarding the uses of music in everyday life (see Related Work section for an excerpt). A literature survey should give me an overview of relevant patterns and be able to deduct which of these can be found in digital listening histories.

(1b) Second, a data mining and statistical approach with cumulated listening histories should allow me to

---

<sup>1</sup> <http://www.last.fm>



Figure 1 - Song- and listening session-centric visualizations *Tangle* (top) and *Strings* (bottom) [1].

see the relevant aspects of single histories and determine what makes an individual history relevant by putting it in contrast to others. Quickly extracting these aspects from a given history can improve personalized interfaces. As an additional benefit, this could also lead to new insights into what types of listeners exist. (1c) Finally, listening histories reflect decisions and the musical taste of a person and as this information is only partially contained in the data I need input from the users themselves. Listening histories with thousands of entries can become quite complex, so I must give the users the tools to understand them: I develop and implement visualization prototypes that shed light on different aspects of listening histories and help uncover patterns that we were not aware of before. Subsequent user studies in the lab and in the wild can uncover those.

These visualizations also constitute one answer to my second research question: Giving users the tools to learn about themselves has benefits on its own. (2a) Additionally, I want to perform another literature survey on user tasks and interfaces surrounding music. Its outcome, the design space of different types of music tasks and applications, should provide the basis for mapping the patterns found in listening histories to these applications and improving a user's experience. (2b) Finally, I will develop concepts and prototypes showing different ways to integrate this data. If there is sufficient time left, I would like to evaluate these prototypes as well.

### Completed and proposed research

At the moment, both literature surveys on listening behavior and music-related applications (1a and 2a) are close to finished and I am currently working on

structuring the literature and deriving the design space. I also (1b) analyzed a set of 5.000 long-term listening histories using PCA to uncover the most salient features of listening behavior. The results especially showed the impact of seasons which was not observed before as this is the first study with listening histories spanning more than two years. Also, character traits that influence listening such as curiosity were uncovered.

The main part of my Ph.D. projects concern creating visualization prototypes and making them available (1c, 2b):

As a first experiment, *Strings & Tangle* [1] (see figure 1) provides song- and session-centric views on small listening histories. The two visualizations gave an initial impression of this data type and have been evaluated informally. The most sophisticated prototype so far, *LastHistory* [4] (see figure 2), allows visualizing listening histories of arbitrary sizes in a time-centric way. The two-dimensional graph of songs (horizontal: days, vertical: hours) shows daily and long-term patterns. Analysis tasks such as searching and filtering, and interactively finding song repetitions and sequences are possible. Additionally, contextual data in the form of photos and calendar entries can be overlaid as memory triggers that ease connecting the abstract song titles to one's personal history for reminiscing. We qualitatively evaluated *LastHistory* with four users in the lab and made it available on the internet. It was downloaded by several thousand people and also discussed by several technology blogs. Informal feedback and an attached online questionnaire allowed us to learn how people used the application and about the patterns they found. With a similar approach, i.e., making the prototype available on the web, I currently explore the connections between multiple listening

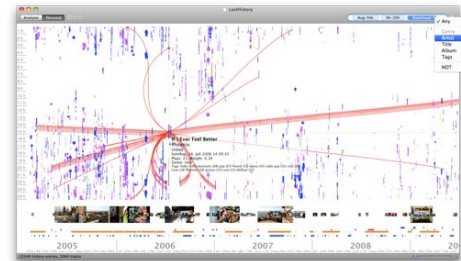


Figure 2 - Visualizing a listening history, photos and calendar entries in *LastHistory* [3].

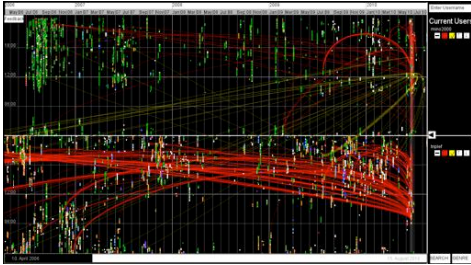


Figure 3 - Two concurrent listening histories visualized for comparison in *LastLoop*.

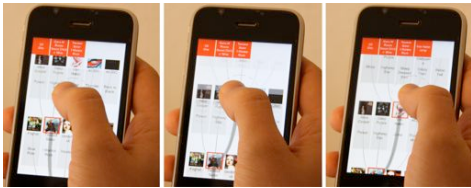


Figure 4 - *Rush*: Interactive, intelligent playlist-building on a mobile device [2].

histories: *LastLoop* (see figure 3) gives users the chance to compare their listening behaviors with those of friends and family. In addition to answering a questionnaire, they can also take screenshots and explain what they found at that moment. The feedback from the prototypes showed that people indeed want alternative ways to access their listening histories. Another common task for music listeners is creating playlists (also 2b). One concrete prototype that I developed in this regard is *rush* [2] (see figure 4), an interaction technique for interactive recommendation on mobile devices. For a given seed song the user is getting a selection of five related songs. After choosing one, another set of related songs becomes available and so on. While the first prototype did not integrate listening histories, I want to use them to make sure that the user knows the suggested songs. Also, personalized recommendations could improve the experience. The original *rush* was evaluated in a lab setting, so I plan to make it available on Apple's AppStore to reach a large number of users. I also want to build or at least design an additional prototype that integrates listening histories with the (currently) tedious task of organizing one's music collection and evaluate the impact. Finishing these additional adaptations and prototypes should take no longer than half a year.

### Expected contributions

The outcomes I expect from my thesis for the HCI community are a deeper understanding about the ways people listen to music, which can inform all types of music-related user interfaces and recommender

systems. Also, visualization researchers can learn how this specific type of data, but also personal or lifelogging data in general can be visualized and what influence a user's memories have in the process. And finally, listeners gain novel ways to work with their music.

### References

- [1] Baur, D. and Butz, A. Pulling Strings from a Tangle: Visualizing a Personal Music Listening History. *Proc. IUI* (2009), 439-444.
- [2] Baur, D., Boring, S., and Butz, A. Rush: Repeated Recommendations on Mobile Devices. *Proc. IUI* (2010), 91-100.
- [3] Baur, D., Seiffert, F., Sedlmair, M., and Boring, S. The Streams of Our Lives : Visualizing Listening Histories in Context. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1119-1128.
- [4] Herrera, P., Resa, Z., and Sordo, M. Rocking around the clock eight days a week: an exploration of temporal patterns of music listening. *Proc. WOMRAD Workshop (in conjunction with ACM RecSys)* (2010), 7-10.
- [5] Pousman, Z., Stasko, J., and Mateas, M. Casual Information Visualization: Depictions of Data in Everyday Life. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1145-1152.
- [6] Rentfrow, P.J. and Gosling, S.D. The do re mi's of everyday life: The structure and personality correlates of music preferences. *Journal of Personality and Social Psychology* 84, 6 (2003), 1236-1256.
- [7] Sellen, A. and Whittaker, S. Beyond total capture: a constructive critique of lifelogging. *Communications of the ACM* 53, 5 (2010), 70-77.