

arcs.fm - A Backdrop Visualization for Music Talk

Dominikus Baur¹, Andreas Butz² and Sheelagh Carpendale¹

¹Department of Computer Science, University of Calgary, Calgary, Alberta, Canada

²Institut für Informatik, Ludwig-Maximilians-Universität, Munich, Germany

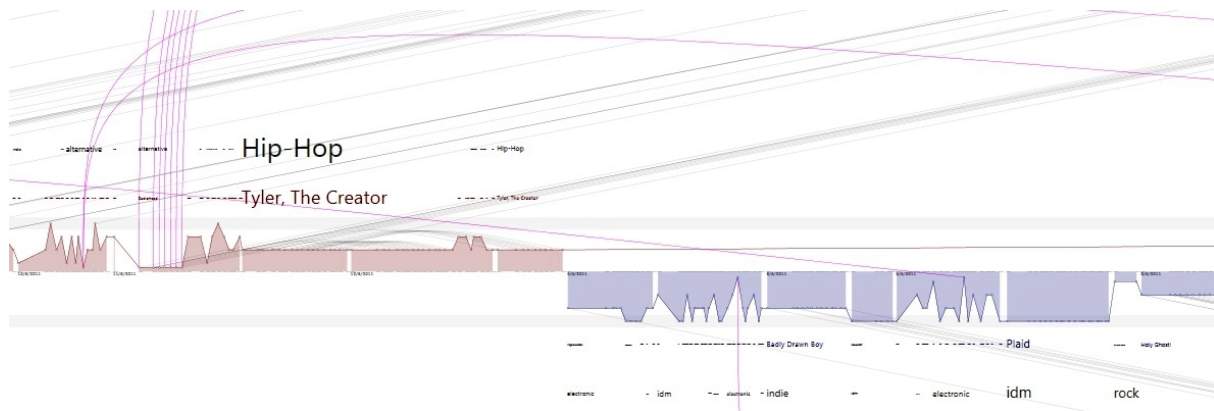


Figure 1: *arcs.fm*, a backdrop visualization for discussions about music, compares two personal music listening histories.

Abstract

Visualizations usually completely capture our attention or disappear into the ambient background. In this paper we explore a middle ground, with visualizations that are not constantly in the center of attention and support a main conversational task without distracting from it. Such backdrop visualizations work for look-up and analytical tasks without getting in the way of the conversation, but can also be used actively.

To illustrate this concept we describe *arcs.fm*, a case study for music talk, as an exemplary backdrop visualization. Music fulfills an important function for identity construction: we define ourselves by what music we listen to and we like to compare our musical taste with friends and family. The shift towards digital music allows us to meticulously keep track of all songs we have listened to and to have access to this data to augment our memories. Here we explore integrating visualizations of automatically collected listening histories with the explanations and discussion that develop in personal music talk. *arcs.fm* is a visualization system that supports this music talk by comparing two listening histories visually. *arcs.fm* stays in the background while enabling look-up and enriching peoples' conversation when needed.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Information Interfaces and Presentation]: User Interfaces—Graphical User Interfaces (GUI)

1. Introduction

In this paper, we present the concept of *backdrop visualizations* – visualizations that support conversations without interrupting them. Our term comes from photography where backdrops form a frame without distracting from the

main motif. Similarly, backdrop visualizations support the main task of discussion without actively drawing attention to them. This does not mean, however, that backdrop visualizations completely disappear into the ambient environment. They also work as look-up tools that sometimes become the

center of attention. This means that such visualizations have to fulfill two roles – a passive one: as non-distracting but inspiring backgrounds, and an active role: as effective data look-up tools, and that they switch between those states at random intervals depending on the conversation.

arcs.fm is one specific example of a backdrop visualization for the domain of music talk. The visualization shows two personal music listening histories as a background for the conversation. From time to time, the conversational partners might use it for two active tasks that are central for backdrop visualizations: (1) for quick look-ups of longer and shorter periods of time (e.g., to find the name of a song), or (2) for exploration and highlighting interesting musical correlations which neither conversational partner is aware of (e.g., a shared liking for a specific artist).

Music is an appropriate topic for this exploration because it plays an important role not only in upholding one's self-image but also in determining one's perception by peers: adolescents frequently engage in discussions about music [RMO09] and even later in life, taste in music still retains an important function for identity management. A music listening history is an accurate reflection of this taste in music, but framing it and providing commentary is important. The power of music talk (we use the term following Frohlich et al.'s notion of *photo-talk* for similar discussions about photos [FKP*02]), comes from the freedom to emphasize personally important aspects, providing a personal perspective on events, being able to quickly switch facets and the general to-and-fro between participants.

While people shine with these verbal activities, their memory is not flawless: even when discussing recent musical favorites they might forget about some of them and providing an accurate overall description of how one's taste in music developed is difficult. In addition to this frailty of memory, music talk only uncovers interesting similarities such as a shared favorite artist or knowledge of an uncommon genre by accident. Using visualization in concert with automatically collected music listening histories can enhance these tasks.

2. Related Work

The recent trend of lifelogging [LDF10] has led to more and more research on the potential of visualizing such personal data. In the domain of music, existing visualizations usually focus on single listening histories: Byron and Wattenberg's streamgraphs [BW08] give an abstract overview, while Baur et al.'s visualizations focus on song sequences [BB09] or adding contextual information for reminiscing tasks [BSSB10]. Comparing multiple histories is more common in non-scientific contexts: last.fm, a web-service that collects listening histories, provide a timeline-based visualization with a focus on song frequency (*Scrobbling Timeline*, <http://playground.last.fm/demo/timeline>) and

Last.fm Explorer (http://alex.turnlav.net/last_fm_explorer/) provides a streamgraph for two histories and even allows drilling down to the level of songs (with corresponding visual clutter). Chen et al. present in a workshop paper two complementary visualizations for comparing two listening histories on different levels of abstraction [CBB10] that focus on tag overlap and time, respectively. Our own work is based on improving on their time-centric second visualization, LoomFM.

Having discussions form over visualizations commonly happens in collaborative visualization scenarios (for an overview, see [IES*11]). Collaborative visualization, however, puts the analysis and gaining insight tasks first: Conversation is only used as part of the sensemaking process. In contrast, backdrop visualizations include aspects of analysis but are much less formal and only supportive – the discussion could also take place without them. DiMicco et al. also presented abstract background visualizations for displaying participant involvement in meeting scenarios [DB07], but these work more as feedback than basis for the conversation. In this regard, backdrop visualizations such as *arcs.fm* have much in common with casual [PSM07] and ambient visualizations, but with a higher involvement – conversationalists can make use of the visualization but are not necessarily relying on it.

3. *arcs.fm*

In this section we describe the characteristics of our visualization and the rationale behind its design.

3.1. Data and meta-dataset

Last.fm (<http://www.last.fm>) is a web-service that collects listening histories with their client software. They expose this data via an API that not only provides lists of songs and timestamps for each profile, but also additional meta-data, such as genre, artist images, etc. Listening history data from last.fm comes as pairs of timestamps and song IDs. For each song we extracted the artist's name and image. Additionally, we retrieved a list of ten similar artists and the ten most popular user-generated keywords. The most popular keyword for an artist on last.fm is coincidentally also almost always the musical genre they belong to [BSSB10].

3.2. Visualization requirements

One aspect that we had to take into account when designing *arcs.fm* was that the visualization would form the backdrop to a conversation. This means that the owners of our histories would usually be deep in conversation. Therefore, looking something up by interaction or finding interesting patterns by understanding should require as minimal an effort as possible. In the best case, *arcs.fm* would only provide cues for learning about interesting aspects and keeping the conversation going. We therefore opted for a visualization approach




Figure 2: Even at the lowest magnification, overarching patterns such as the tendency towards repetition (grey arcs), overlaps in taste (pink arcs) and length of listening sessions (consecutive boxes) are still visible (4000 songs).

with clear visuals and focused on minimal interaction. While backdrop visualizations can potentially provide two different versions of the visualization for their passive and active states, we decided for a uniform approach with arcs.fm to keep distractions and disorientation through switching low.

3.3. Design Rationale

Comparisons require a common ground for the data but also for the conversation. Time is such a unifying dimension, that was both relevant for mapping the data but also for remembering and talking about it. We chose to visualize it as a timeline (for an overview of this technique, see [AMM*07]). As time took up only one of our spatial dimensions, we could use the other for emphasizing the contrast of the histories. We decided that the more similar a song from one history was to the other, the closer it would move towards that other history (*closeness = similarity*). Having only two histories allowed us to place one of them above and the other one below a horizontal central timeline.

In order to provide an overview of the data but also show details-on-demand and at the same time keep the intensity of required interaction as low as possible we opted for a zoomable user interface. The lowest magnification would allow for quick navigation and provide an overview, while it remained possible to zoom in for more detailed information of individual songs. We used the shrinking of items when zooming out for encoding relevance: a label's size and consequent readability is coupled to its relative importance (*size = importance*). Large and important labels would stay readable even while zoomed out, while less expressive ones would still be there but become harder to read.

One aspect of the dataset fruitful for discussions are specific songs that frequent both histories. The higher their frequency, the more expressive and relevant they are for the discussion. Therefore, we used arcs (cf. arc diagrams [Wat02]) for visual highlighting (*arcs = identical songs*).

3.4. Visual mapping and interaction

Figure 1 shows how all songs are aligned to the horizontal timeline and displayed as small circles. The spacing on the timeline is not linear: we used the natural logarithm of the temporal distance between two songs in minutes for their horizontal distance. This creates equal spacing between

songs (their length is usually between 3 and 8 minutes), but drastically shrinks empty (and uninteresting) time segments. In order to make songs easier visible on lower magnification, we added rectangular areas that vertically connect the small song circles to the timeline. Each such area contains all songs from one listening session (i.e., a non-interrupted sequence of songs) and thus highlight these sessions' lengths and frequencies (see Figure 2). As the chances that people listen to music at the same time are slim, listening sessions usually appear alternating (even though parallel sessions are possible). In light of the *label size = importance* concept, both dates and times of day are shown along the timeline, but labels for dates are larger and more readable. Labels are only shown for dates where music was played.

Songs from the first listening history are displayed above the timeline and from the second one below it. Both histories are additionally color-coded (in the figures: red and light blue. Color schemes can be adjusted for color deficiencies). Distance from the timeline and thus from the other history obeys the principle of *closeness = similarity*. While there are various methods of increasing complexity to calculate musical similarity, we wanted to keep it comprehensible. Therefore, we defined five discrete steps from most similar to least similar: (1) *song appears in other history*, (2) *other songs by the same artist appear in other history*, (3) *songs by similar artists appear in other history*, (4) *songs of the same genre (= most popular user-generated keyword) appear in other history*, or (5) *none of those attributes fit*. Two light grey bars above and below the timeline highlight the end points of the similarity scales for each history.

Vertically above/below each song the system displays labels for artists and genres. These labels again follow *size = importance*: If two consecutive songs share artist or genre their two labels are merged into a larger one. This rule leads to labels for prominent artists and genres becoming large enough to be read even on low magnification (see, e.g., Figure 1 – 'Tyler, The Creator' and 'Hip-Hop').

A last aspect of the visualization concerns the repetition of songs. People commonly listen to a song more than once with each instance of it leaving an imprint in the listening history and making them all appear along the timeline. Similar to arc diagrams [Wat02], we drew arcs between instances of such repeating songs. arcs.fm provides two types of arcs: *intra-history arcs* between songs that only appear in a single



Figure 3: Hovering over an arc or song node shows additional information and a listing of all its instances.

history hop from instance to instance, so, e.g., three identical songs in one history leads to two arcs between them. They have grey color and are symmetrical; their height corresponds to one tenth of the distance between song instances. This makes it possible to see how frequently songs are repeated: Low arcs mean that songs are only repeated shortly after discovering them, while high arcs stem from frequent, long-term repetitions. *Inter-history arcs* connect the first instance of a song to all instances in the other person’s history. These are colored in pink and have a skewed curvature that leads to an almost vertical initial shape with a long fall-off (see Figures 1 and 2). These differences in color and shape make them visually distinguishable from intra-history arcs and additionally allow seeing in which history a song appeared first – an important feature for determining who has the ‘bragging rights’ for discovering an artist.

Interaction is intentionally kept simple to prevent it from disrupting the conversation. Zooming and panning are triggered by using the mouse wheel or dragging the canvas, respectively. Hovering over a song circle or arc shows detail information (see Figure 3).

3.5. Implementation

arcs.fm was implemented in C# using the WPF framework. The decision to refrain from using animations allows our system to display two histories with a total of 10,000 songs at interactive frame-rates on an i7 machine with 8GB of RAM on a full-HD monitor (10,000 songs comply with a two-year time-span for an average of 7 songs per day and two histories). In addition, we also created a touch- or pen-enabled version suitable for large wall displays. In this version, the lack of a mouse wheel was countered by making panning one-dimensional along the horizontal dimension and mapping the vertical input dimension to the zoom-level.

4. Discussion

arcs.fm highlights overarching listening behavior (number and distribution of circles = overall song frequency, width of blocks = length of listening sessions, number and height of arcs = repetitive behavior) and also important artists and genres through their label size. Outliers appear in the form

of single song blocks, short isolated listening sessions, but also as song circles without arcs (if songs are frequently repeated otherwise). The overall closeness of the song circles and also the movement of the line connecting all songs of one history (zigzag pattern for frequent disagreements, straight close lines for high similarity) can be used for comparisons. The inter-history arcs show who discovers songs first but also the general overlap between the two histories (number of arcs).

Regarding its efficiency as a backdrop visualization, arcs.fm supports passive exploration through its clear and unintrusive visual mapping and lack of animations. Depending on the current topic of conversation (period of time, overarching listening behaviour, etc.) a different section of the timeline can be brought into focus and discussed without further interaction. Active look-up tasks work with the same zoom- and pan-navigation and hovering for detail information. Relevant intersecting patterns are highlighted through inter-history arcs and larger genre- and artist names. Interaction in arcs.fm is so basic that conversational partners necessarily return to their conversation quickly after adjusting the timeframe or showing detail information for a song. More complex interactions for filtering or searching would have introduced modes and were deemed too distracting.

Music talk is not the only conversational task that could be supported by visualization: Backdrop visualizations could also enrich photo talk [FKP*02], discussions about sports (showing information about teams, leagues and players) or politics (with relevant statistics). Similar to arcs.fm, such visualizations should enable looking up facts, but also find and highlight relevant patterns in the data. As using visualizations as conversational backdrops is a new concept, mapping design recommendations to scenarios makes for required future work. Our experience with arcs.fm indicates that the following directions are worth exploring: (1) reducing interaction and thus distraction, (2) reducing ambiguities in the visualization and avoiding modes and corresponding switches and (3) avoid distracting peripheral animations.

5. Conclusion

In this paper, we presented the concept of using backdrop visualizations with arcs.fm as one specific example for discussing music taste. arcs.fm supports conversation without interrupting it by providing conversationalists with an access to their past and highlighting differences that they may not be aware of. We plan to investigate further the aspects of backdrop visualizations that become important for other topics/datasets. We also want to explore the concept of distributed scenarios.

Acknowledgements

We thank our reviewers for their feedback and iCORE, NSERC, and SMART Technologies Inc. for funding.

References

- [AMM*07] AIGNER W., MIKSCH S., MULLER W., SCHUMANN H., TOMINSKI C.: Visualizing time-oriented data - A systematic view. *Computers & Graphics* 31, 3 (June 2007), 401–409. [3](#)
- [BB09] BAUR D., BUTZ A.: Pulling strings from a tangle: visualizing a personal music listening history. In *Proc. IUI '09* (2009), pp. 439–444. [2](#)
- [BSSB10] BAUR D., SEIFFERT F., SEDLMAIR M., BORING S.: The Streams of Our Lives: Visualizing Listening Histories in Context. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1119–1128. [2](#)
- [BW08] BYRON L., WATTENBERG M.: Stacked graphs-geometry & aesthetics. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1245–1252. [2](#)
- [CBB10] CHEN Y.-X., BAUR D., BUTZ A.: Gaining Musical Insights: Visualizing Multiple Listening Histories. In *Workshop on Visual Interfaces to the Social and Semantic Web (VISSW2010)* (2010), pp. 1–5. [2](#)
- [DB07] DIMICCO J., BENDER W.: Group reactions to visual feedback tools. *Persuasive Technology* (2007), 1–12. [2](#)
- [FKP*02] FROHLICH D., KUCHINSKY A., PERING C., DON A., ARISS S.: Requirements for photoware. In *Proc. CSCW '02* (2002), ACM, pp. 166–175. [2](#), [4](#)
- [IES*11] ISENBERG P., ELMQVIST N., SCHOLTZ J., CERNEA D., MA K., HAGEN H.: Collaborative visualization: definition, challenges, and research agenda. *Information Visualization* 10, 4 (2011), 310–326. [2](#)
- [LDF10] LI I., DEY A., FORLIZZI J.: A stage-based model of personal informatics systems. In *Proc. CHI '10* (New York, New York, USA, 2010), ACM, pp. 557–566. [2](#)
- [PSM07] POUSMAN Z., STASKO J., MATEAS M.: Casual Information Visualization: Depictions of Data in Everyday Life. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1145–1152. [2](#)
- [RMO09] RENTFROW P. J., McDONALD J. A., OLDMEADOW J. A.: You Are What You Listen To: Young People's Stereotypes about Music Fans. *Group Processes & Intergroup Relations* 12, 3 (Apr. 2009), 329–344. [2](#)
- [Wat02] WATTENBERG M.: Arc diagrams: visualizing structure in strings. In *IEEE Symposium on Information Visualization, 2002. INFOVIS 2002* (2002), IEEE Comput. Soc, pp. 110–116. [3](#)