

# Using Visualization Techniques to Characterize Music Listening Behavior

José Devezas  
Laboratório SAPO/U.Porto  
joseluisdevezas@gmail.com

Sérgio Nunes  
INESC TEC & DEI-FEUP  
ssn@fe.up.pt

## Abstract

In order to better understand how people listen to music, both from an individual and a global perspective, we have created a set of visualizations to help us gain an intuition on the data and formulate our first behavioral hypotheses. We present the set of visualizations we created and the intuition that this process helped us develop in the analysis of a dataset of Last.fm user profiles. Specifically, we were able to notice that, sometimes, users tend to associate the playback of specific artists with a given weekday. The same happens when avoiding particular artists. Since there was no visible clear pattern that would generalize this behavior, we devised a heuristic to identify preferred and avoided weekdays for a given user and artist. In this work, we focus on describing the iterative process we followed to carry our exploratory analysis supported on data visualization.

## 1 Introduction

In addition to gaining knowledge about the data, visualization also provides a good way of communicating the progress of research work to colleagues and the general public, opening space for discussion and new ideas. Our main goal is to show that this, perhaps less conventional, methodology can be quite rewarding, specially when dealing with the lack of a beforehand hypothesis and the necessity of continuously sharing results.

The case study we present here is based on the characterization of music listening behavior. The goal was to understand whether or not there are differences or similarities between individual and global behavior, and eventually extend to the lev-

els in between. We departed from this generic goal, without having a specific hypothesis in mind, simply because we assume a close-to-zero knowledge about the data.

## 2 Reference Work

The work we present here might be considered visual data mining. Keim [5] provides a description on information visualization and visual data mining, with an in-depth classification of available techniques, according to three axis: data to be visualized (e.g. one/two/multi-dimensional, text, hierarchical), visualization technique (e.g. stacked, 2D/3D), and interaction and distortion technique (e.g. filtering, zoom). This is an extremely good reference for anyone looking for the ideal plotting technique for a particular dataset.

For our specific case, we started with a type of visualization that, while more recent and not covered by the previous article, fits perfectly into their classification scheme under the “three-dimensional”, “stacked” and “interactive linking and brushing” classes. This visualization is the streamgraph [2], which depicts the evolution of categorical data over time — in our case, the evolution of play counts, for each artist, over the course of several months. We did a web-based implementation of this and other visualizations, based on D3.js [1], which already supports many of the atomic operations that ease the creation process through the manipulation of a DOM (typically HTML or SVG).

Herrera et al. [4] present a study, based on circular statistics, where they also found a preferred listening day for certain user and artist combinations. This supports the work we present here and the hypothesis we formed through our visual data mining

process. The difference between the two works is that we focus on the methodology of forming an hypothesis on unknown data, while they started at a later phase of the process, from an already formulated hypothesis, which they tested and proved.

Sarmiento et al. [7] describe a music visualization based on a similarity network of artists, called RAMA (Relational Artists MAp). According to Keim [5], this visualization would fit the “graphs”, “zoom”/“filtering”, and “standard 2D” classes. RAMA is an example of a visualization as a research communication by-product that is delivered as a production-ready prototype. Next, we present the set of music visualizations we created, and the exploration process that led us to our hypothesis.

### 3 Last.fm Dataset – 1K Users

The analysis we present here is based on the *Last.fm Dataset – 1K Users*<sup>1</sup>, provided by Òscar Celma [3, Chapter 3], and extracted from Last.fm<sup>2</sup>, a popular music social network.

The dataset contains scrobble records for 992 users, where each record shows the title and the artist for a track that a specific user played at a given time, ranging from January 27, 2008, to May 4, 2009. There is information for 19,150,868 scrobbles, aggregating 176,948 different artists. This dataset is particularly interesting, as it contains timestamps, enabling the study of taste change and evolving music listening behavior.

## 4 Visualizations

In this section, we illustrate the whole explorative process, including the constraints of the analysis, the ideal visualization techniques to explore our data, the actual observation phase, and the formulation of an hypothesis based on interesting peaks, depressions or other patterns found in the charts.

### 4.1 Setting Constraints

The first step to start exploring the dataset was to decide on which dimensions we should focus in. In our case, since we wanted to understand how

<sup>2</sup><http://www.last.fm>

users listened to music, so that we could then work on prediction algorithms to recommend new music, we were specially interested in centering around the user. We wanted to look at the total playback, as well as the per artist playback, over time, aggregated by different units of time. Characterizing user listening behavior over time has the potential to enlighten us on the process of music discovery (or rediscovery), as well as on the evolution of musical taste. We chose the artist as the “music unit” since we believe it is specific enough, yet still contains enough information to describe listening behavior (genre might have been too broad and track might have been too specific). While we wanted to obtain visualizations per user, we also wanted to understand the global behavior of the users, together as a unit, so that we could comprehend where each individual stood in the whole. Thus, we also analyzed the aggregation of the playbacks for all users, trying to find different behaviors, namely dependent on the gender, but also on the country of origin.

### 4.2 Deciding on Techniques


First of all, we knew that we would mostly focus on studying the evolution of music playback, so we had to depict behavior over time. For this, we mostly used two types of charts: streamgraphs and bar charts. Streamgraphs are a type of stacked area chart, where the x-axis has a fluctuating zero, optimized to facilitate the reading of data with large peaks without impacting the readability of lower values as strongly as traditional stacked area charts.

For the study of the global music listening behavior, we plotted a set of static bar charts aggregating play count data by: hour of the day, distinguishing between gender using stacking and different colors; and hour of the day, distinguishing between country of origin using faceting, and different colors to depict different types of music behaviors (discovered using a clustering algorithm).

### 4.3 Looking at the Data

In this section, we describe the visualizations we created<sup>3</sup>, how they interact with each other to sup-

<sup>3</sup><http://josedevizas.com/#!/individual-listening-behavior> and <http://josedevizas.com/#!/global-listening-behavior>



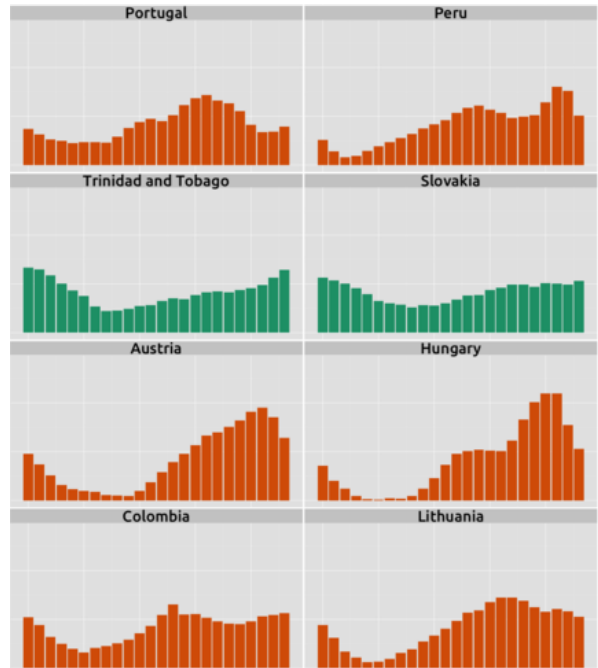
**Figure 1:** Dodged bar chart depicting the hourly distribution of global play counts, from 0 to 23 hours, distinguishing between Male (blue), Female (pink) and N/A (black).

port the exploration process, and what they can show us about the data.

#### 4.3.1 Global Music Listening Behavior

To get an overview of the global music listening behavior, we focused on hourly patterns. First, we plotted the hourly play count distribution, distinguishing between male and female in order to look for a possible gender bias. Figure 1 depicts this behavior, showing a population with more male than female individuals, both presenting a similar music listening pattern, highly correlated with sleep hours.

We then plotted the hourly play count distribution, aggregating by country instead of gender, in order to look for a possible country bias. While in the previous plot male and female clearly showed a similar hourly behavior, in this case, since we were using a faceted plot containing well over two different facets, we used a clustering algorithm to aggregate countries with a similar hourly music listening behavior. Figure 2 shows a partial view of this faceted bar plot with different colors for different behaviors. We used the  $k$ -means algorithm, with  $k = 2$  as the best value for  $k$ , according to the silhouette method [6]. We found two different behaviors, as depicted in Figure 3. Figure 3a depicts a behavior where people listen to less music during work hours, approximately between 6h and 18h, with a minimum listening activity at 10h. Figure 3b depicts a behavior where people listen to less music during sleep hours, approximately between 1h and 9h, with a minimum listening activity at 6h. While this could have been an interesting finding, the two  $k$ -means centers actually illustrated a



**Figure 2:** Partial view of the hourly distribution of global play counts, from 0 to 23 hours, aggregating by country using faceted plotting. Different colors represent distinct behaviors discovered using the  $k$ -means clustering algorithm by setting  $k = 2$  according to the silhouette method.

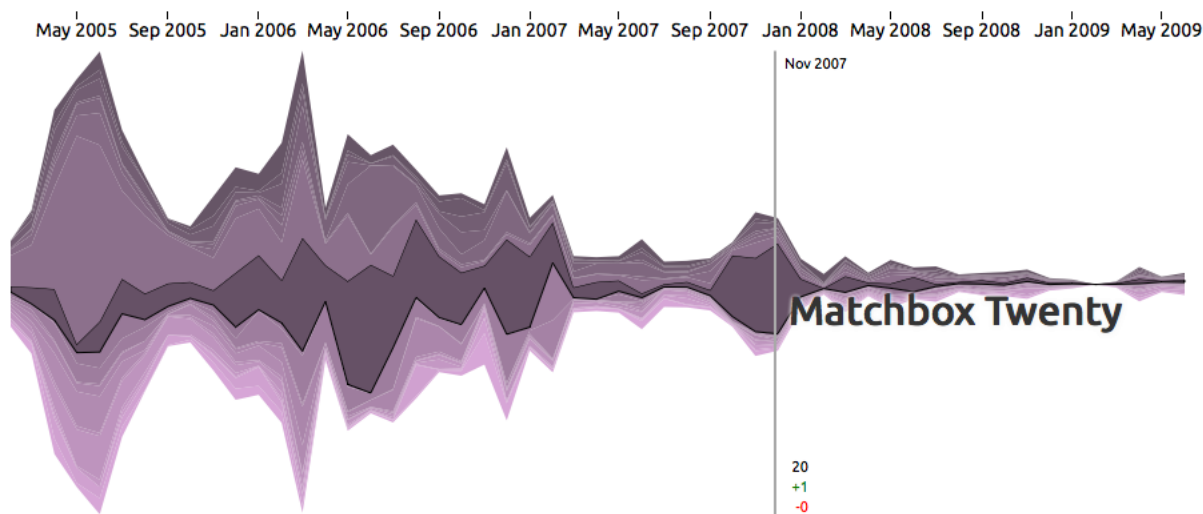


(a) People listen to less music during work hours (~6h-18h).

(b) People listen to less music during sleep hours (~1h-9h).

**Figure 3:** Two different behaviors depicted by the  $k = 2$  centers of the  $k$ -means algorithm.

similar behavior with only a slight temporal shift. This might have been caused by our attempt to convert timestamps to UTC based on the average timezone of each country, since the dataset provided the country and the local time, but not the timezone. While the clustering methodology and its translation to a visual feature is still worth mentioning, we did not assign any importance to this finding, disregarding these results.




**Figure 4:** Streamgraph depicting the artist play count for a given user. We can see the “Matchbox Twenty” highlighted through mouse over and showing: in black, the total number of distinct artists played in November 2007; in green, the number of artists added to the playlist compared to the previous month; and, in red, the number of artists added to the playlist compared to the previous month.

### 4.3.2 Individual Music Listening Behavior

Figure 4 illustrates the first graph we created to analyze the user-centric behavior. This streamgraph shows the play count, aggregated by year and month, for a given user, selectable from a drop-down box. Different artists are shown with different tonalities and, on mouse over, the name of the artist corresponding to the highlighted portion of the graph is displayed (e.g. “Matchbox Twenty”). While moving the mouse, we provide a vertical guide to help determine the position in the time scale (e.g. “Nov 2007”), as well as to keep track of the total number of artists at a given time (e.g. 20), and of the number of new (e.g. +1) and dropped (e.g. -0) artists when compared to the previous month. On mouse click, the associated bar charts are updated, showing data for the selected user and the clicked artist.


Figure 5 essentially translates the selected portion of the streamgraph, corresponding to an artist, into a bar chart. This view enables an easier perception of the quantity of monthly playbacks for a given user and artist. Nearby the chart, we also provide a drop-down box for quickly switching the displayed artist. Exploring different artists and their playback over time helped us gain an intuition on how a particular user listened to music or even



**Figure 6:** Bar chart depicting the total number of play counts aggregated by weekdays for “user\_000298” and “Matchbox Twenty”.

if different users followed similar patterns. We also found this visualization particularly relevant in understanding the timespan of data available for each user and the quantity of playbacks for each artist, which supported the exploration process by serving as a sort of relevance meter to help decide whether or not a particular behavior of interest could be supported by the data.

When an artist is selected, a bar chart depicting the daily behavior over a weekly period is also



**Figure 5:** Bar chart depicting the total number of play counts over the months for “user\_000298” and “Matchbox Twenty”.

updated to show data for that particular user and artist. Figure 6 illustrates an instance of this visualization with the mouse hovering the Wednesday bar and showing the total number of playbacks for that day.

#### 4.4 Formulating an Hypothesis

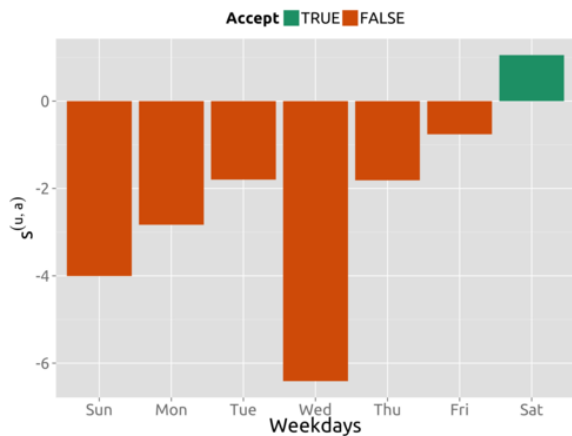
While the graphics we created provided a good visual translation of our data, we did not detect any interesting pattern until we started analyzing the weekly behavior of each user. At this point, we started noticing that a large number of users seemed to prefer to listen to some artists at a particular weekday, having one, and sometimes two, larger than average bars. We also noticed that, in some cases, the reverse would happen, as there was one or two days when a user specially avoided listening to an artist. This pattern seemed coherent with natural human behavior and the weekly cycle of work life. However, we expected perhaps a more weekend focused change of behavior, which didn’t seem clear from analyzing the visualizations, given many users preferred to listen to some artists in the middle of the week, for instance. Thus, the first hypothesis we formulated by looking at the data was whether or not users had any preferred and avoided days to listen to a given artist. Other questions also started to arise. Is the behavior similar across different artists? And across different users? If so, can it be explained by the influence of the weekend? Or does the person work at weekends? Is a particular musical genre associated with a particular day of the week? For individual users, or for all of them?

As you can see, even if we assume to have zero knowledge about a dataset, we can use visualization to guide our curiosity towards establishing an initial hypothesis. And while visualizations can be quite illustrative, many plots might not add much beyond confirming the expected behavior. On the other hand, it makes sense to keep pushing the analysis until we are able to understand as many dimensions of our data as possible. This is a simple methodology that many senior researchers follow instinctively, but that might be of special interest to rookie researchers in order to help them prepare their minds towards the scientific method. We point out that, while apparently, a scientist starts with an hypothesis, that is not always the case, and an hypothesis must be first found. One way to do so is to analytically study the data, but also to use visualizations to interpret it.

## 5 Weekday Artist Bias

We used the visual intuitions as a cue to propose the following hypothesis: a user has preferred and avoided weekdays for the playback of certain artists, where preference implies a diminished interest in the remainder artists during the day. We developed a heuristic to identify such artists and weekdays, and linked this new information to the visualizations as a way to feed the exploration process forward. While we still consider this to be preliminary work, we describe this heuristic next.

Assume a vector  $p^{(u)}$  where each element corresponds to the total play count, per weekday from



**Figure 7:** Weekday preference score  $s^{(u,a)}$  for  $u = \text{user\_000196}$  and  $a = \text{Lamb}$ .

Sunday to Saturday, for all artists of a given user  $u$ . We normalize this vector by dividing each of the seven elements by the maximum play count of the vector. Assume a vector  $p^{(u,a)}$  where each element corresponds to the total play count, per weekday from Sunday to Saturday, for a given user  $u$  and a given artist  $a$ . Similarly, we normalize this vector by dividing by the maximum play count of the weekday totals. Both  $p^{(u)}$  and  $p^{(u,a)}$  are values between 0 and 1, representing the daily fraction of music played over a week.

We then calculate  $p'^{(u,a)} = p^{(u,a)} \circ \hat{p}^{(u)}$ , where  $\circ$  denotes the element-wise product and  $\hat{\cdot}$  denotes the element-wise inversion of a vector, i.e.  $(\hat{p})_i = p_i^{-1}$ . The elements of  $p'^{(u,a)}$  thus work as a metric where larger values indicate artists that are played more on weekdays with fewer global plays. We assumed this metric to be a local indicator of preference. However, to really decide whether or not a particular user has a preference for an artist at a given weekday, we must compare the  $p'^{(u,a)}$  values for the artist with the ones for the remaining artists, deciding on the “closeness” of each value to all others.

Thus, given a user  $u$  and an artist  $a$ , we uncover any preferred playback weekdays by calculating, for each element of  $p'^{(u,a)}$ , the sum of the differences between every other element and the current element multiplied by a weight  $w$ , that should be set to higher values in order to increase preference strength — for our study, we used  $w = 1.5$ ; decreasing it to  $w = 1$  would frequently result in

three or more preferred weekdays; increasing it to  $w = 2$  would frequently result in no preferred days, even when there was a clear bias towards playing the artist at a given weekday. The final result is a vector of scores  $s^{(u,a)}$  that can be used to identify possible preferred playback weekdays for a user  $u$  and an artist  $a$ . Figure 7 depicts the score values calculated to test the hypothesis that “Lamb” has a preferred playback weekday for “user\_000196”. This score should be positive for days of the week where the user is more inclined towards listening this artist. In this case, “Lamb” seems to be a Saturday artist for “user\_000196”.

## 6 Conclusions

We have described a methodology supported on visualization to explore an unknown music dataset, departing from zero-knowledge about the data and arriving at one or more hypotheses. We have shown that, while some charts depict the expected behavior, simply confirming the norm, other charts can lead to quite unexpected and interesting hypotheses. While the ideas we present here lack a high degree of novelty, we aimed at better describing a process that, because of being common knowledge for experienced researchers, is not well documented and available as a resource for beginning researchers.

### 6.1 Future Work

As future work, we would like to point out two possible paths. The first is regarding the continued exploration of the hypothesis we discovered, further studying the user preference or avoidance of particular artists at a given weekday. The second is regarding the management of relevant visualizations as they are added to the analysis, specially in respect to how they interact with each other, and which views are presented in order to keep the visualization dashboard clutter-free and useful in the exploratory process.

## Acknowledgments

This work has been partially financed by Laboratório SAPO/U.Porto under a research grant for the Juggle project.

## References

- [1] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D3: Data-Driven Documents. *IEEE transactions on visualization and computer graphics*, 17(12):2301–9, December 2011.
- [2] Lee Byron and Martin Wattenberg. Stacked Graphs - Geometry & Aesthetics. *IEEE transactions on visualization and computer graphics*, 14(6):1245–52, 2008.
- [3] Òscar Celma. *Music recommendation and discovery in the long tail*. PhD Thesis, Universitat Pompeu Fabra, 2009.
- [4] Perfecto Herrera, Z. Resa, and Mohamed Sordo. Rocking around the clock eight days a week: an exploration of temporal patterns of music listening. In *Proceedings of the 1st Workshop On Music Recommendation and Discovery (WOMRAD 2010), held in conjunction with the 4th ACM Conference on Recommender Systems (RecSys 2010)*, pages 7–10, 2010.
- [5] Daniel Keim. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics*, 5(2):123–8, 2002.
- [6] Peter J. Rousseeuw. Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987.
- [7] Luís Sarmiento, Fabien Gouyon, Bruno G Gosta, and Eugénio Oliveira. Visualizing networks of music artists with RAMA. In *5th International Conference on Web Information Systems and Technologies (WEBIST 2009)*, 2009.