A GRAPHICAL REPRESENTATION OF DIGITAL MUSIC LIBRARIES USING RELATIONAL AND ABSOLUTE DATA

TO REDISCOVER YOUR COLLECTION

Lisa Dalhuijsen

Supervisor: Edwin van der Heide

MSc Media Technology
Leiden University, The Netherlands
2011

ABSTRACT

This paper proposes a visualization system for digital music collections. The system uses both relational data (e.g. similar artists) and absolute data (e.g. song tempo and genre). All of the data is readily available for all popular music and require no manual work from the user. Artists, albums and songs are represented as nodes in a network. A force based graph determines the positions of the nodes on screen. User test results show that the new system helps the users to find the music they want to hear. The visualization provides a truly new view on the information that can't be easily conveyed in text. It provides a good overview of the music at a glance and offers new insights into the collection.

Keywords: data visualisation, digital music, interface design, force based graph

INTRODUCTION

I am both a heavy computer user and music lover. I listen to music on my computer daily - if not all day - and have quite a big collection of MP3 files. I started thinking about this project out of frustration with existing software for creating music playlists. While our collections are ever growing, it is increasingly hard to keep them well-organized. And even in an organized library it might be very difficult to find a specific song. Today's well-known music players lack a certain degree of flexibility and creativity. And when software is being used on a daily basis, even minor inconveniences can lead to great dissatisfaction.

This paper describes an investigation into new ways of how digital music libraries could look and behave, and presents a prototype solution.

THE PROBLEM

Firstly, most existing music libraries like iTunes and Winamp fully rely on a text-based approach. That means that users with extensive libraries will have to wrestle through huge lists in order to find something. Of course the programs offer search functions. So if you (partially) remember the name of the song, album or artist, you're fine. But if you don't, such search functions are of little help.

Another frustration is the way in which digital music is classified. An MP3 file has so-called ID3 tags containing text, that can be edited by anyone. These tags contain information about the music, like the name of the artist and the year in which the music was produced. It can also contain the genre of the piece of music, like *rock* or *classical*. Most music, however, does not fit into one category but is rather a mixture of different styles. With existing music browser software, it is not quite straightforward to find funky jazz or electronic reggae. And finding laid-back music for a sunday afternoon (regardless of the genre) is even more difficult.

The final main problem is that in available music programs it is hard to get a good overview of the music in the collection. Usually, one can browse by artist, album, etc., resulting again in alphabetic lists of text, somewhere containing the music you want to hear.

RESEARCH QUESTION

Based on the fact that humans are good at recognizing visual properties and patterns, my hypothesis is that a smart way of visualizing a music collection will reveal much more information to the user in just a glance, than the all familiar text representations. This leads to the following research question:

› *Which musical properties, which translation of those properties, and what kind of visualization system help users to find the music they want to hear?*

The scope of musical properties should be limited to properties that are generally and readily available for all music. For example, lyrics will not be included because they are limited to music with vocals. Automatic genre and mood classification are not considered because it is not yet possible to determine those automatically with sufficient probability [MIREX 2009 and MIREX 2010].

Because most people organize their music collections by artist/album/song [Vignoli 2004], a hierarchy like that should be included. On the other hand, nowadays people generally buy more individual songs (for example from iTunes) than whole albums, so it might be that a song-level interface is becoming more appropriate. The new system should include both ways of grouping in order to test which one is of better use.

In order to evaluate the resulting system, the following questions will be asked:

› *Can people find the exact artist/album/songs they are looking for?*
› *Can people find the kind of songs they would like to hear, for example songs that fit their mood?*
› *Do people get a better overview of the scope of their music collection?*
› *Are people more likely to discover music in their collection they hadn't heard before?*

MUSICALNODES

In 2007, Lieven van Velthoven and I started working on a solution to the same problem. The system we proposed was called MusicalNodes [Dalhuijsen and Van Velthoven 2010]. In the MusicalNodes project, a physics system coupled to album's genre information allowed the user to spatially order his or her albums on screen and make selections there. MusicalNodes also had a search mechanism and the ability to add custom properties to music.

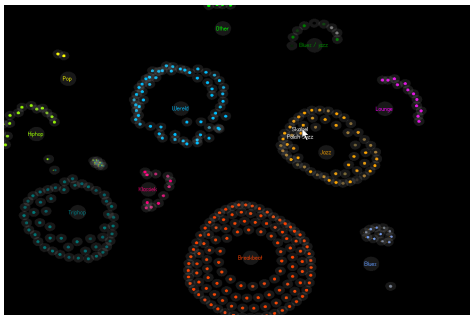

The MusicalNodes project is a step into the right direction. The visualization of music turned out to be a promising approach to organizing music libraries. User tests showed that it offered a useful and playful way to organize music collections. Also, it can provide information that's not easily shown in text, for example music that's hiphop mixed with a little bit of jazz.
But there are still a lot of features to improve.

Firstly, every node represents an album. However, the choice to organise music on album level was quite arbitrary. People don't only organise their music on album level. Sometimes they look for a specific artist, sometimes for a specific song and sometimes for a specific album [Vignoli 2004].

Secondly, MusicalNodes relies on the very effective Gestalt theory for visually grouping objects, but uses it in an inconsistent way. In the system, the album nodes are attracted to the genre node(s) it belongs to. But using proximity to show a relationship also suggests that two neighbouring nodes are musically similar, which is not the case in MusicalNodes. This works counterintuitive to the user.

Another disadvantage of MusicalNodes is that while it is very flexible, it requires a lot of its users since it is so dependent on manual classification of music in the library. The program gets basic information from ID3 tags. But if you want to do something extra (like assigning multiple genres to a song or adding a value for a custom property) you have to do it manually. You can do without custom properties, but then MusicalNodes is actually nothing more than a visualization of the same information iTunes shows you: artist name, album title, album year, song title and genre.

RELATED WORK

Several projects have been done in this field by others.

PlaySOM

PlaySOM [Dittenbach et al. 2005] uses sound signal analysis to extract audio features from songs. That data is used to train a self-organizing map, in order to automatically create a two-dimensional map of the songs. When the user zooms in, more detailed information about the songs becomes visible. The layout of the songs on screen, combined with the heatmap-like display, seems to very intuitive. Unfortunately this system can only be used on song level.

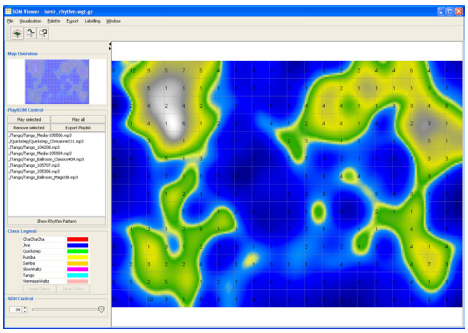

There are quite a few projects like this coming from the research field of *music information retrieval*, for example **Margrid** [Julia and Sorda 2009], **SOMeJB** [Rauber et al. 2002] and **NepTune** [Knees et al. 2007]. They all use self-organizing maps to position songs in a 2- or 3-dimensional space, based on audio features. The user interfaces are generally not well developed, because they are meant to be research projects rather than ready-to-use products.

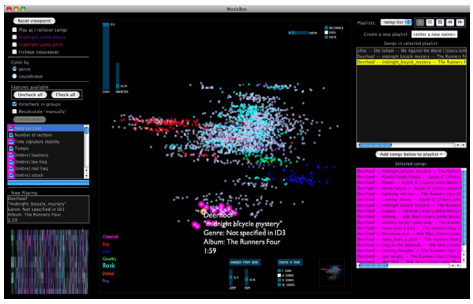

MusicBox

MusicBox [Lillie 2008] shows songs in a scatterplot. For every song a whole list of variables can de included to determine the position of the song, ranging from low-level audio features to the number of tags Last.fm users assigned to that song. Principal component analysis is used to reduce those variables to two dimensions. MusicBox offers a lot of flexibility about which properties of the music are used to create the visualization. Unfortunately, again it's only on song level and not easily scalable to artist or album level.

The Artist Map

The artist map [Vignoli et al. 2004] calculates similarity between artists from audio features of the artists' songs. The artists are then displayed as nodes in a force directed graph. Futhermore, artist nodes can be attracted by *property magnets* that represent genre, *mood* or year of production. The system manages to combine different kinds of information nicely. Low level spectral properties are used but hidden, while at the same time properties that the user might be looking for (such as tempo) are visible. Unfortunately the system is not completely realistic for everyday use: the genres are manually labeled and the *mood* of music is taken from MoodLogic database (which contains a very limited set of music and hasn't been maintained since 2003).

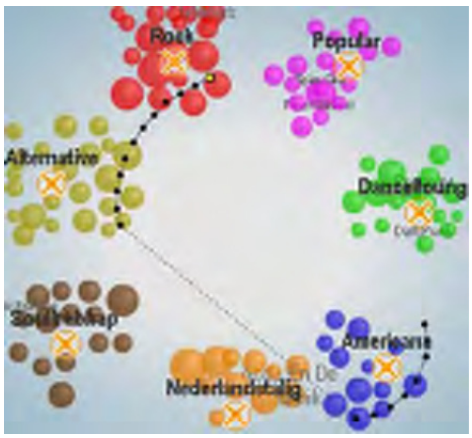

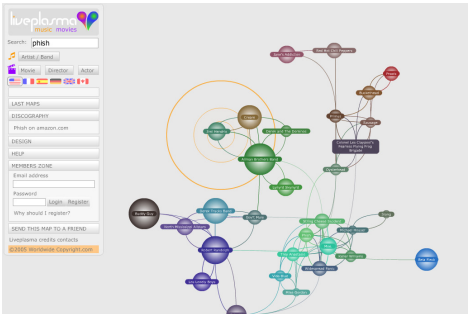

Liveplasma

Liveplasma (www.liveplasma.com) is a visual music discovery tool that asks the user to enter the name of an artist they like. Liveplasma then shows related artists, based on data from Amazon (*"people who bought this album also bought ..."*). Basic information about the artists and their releases is shown.

Quite a lot of projects like this exist online, for example **Tuneglue** (audiomap.tuneglue.net), **Music Map** (www.music-map.com) and **Discovr** (www.discovrmusic.com). They are simple programs that use spring-like networks to show recommendations from online databases like Amazon or Last.fm (*"people who listen to this artist also listen to ..."*). All of them are on artist-level and allow the user to discover connections between artists.

OUTLINE

The projects described above use either properties of the music itself (like timbre or genre) or relationships between different pieces of music (*"people who listen to this artist also listen to ..."*). They never include a combination of those two, which could be double as good. So the new system should include both kinds of information.

The system should be suitable for music collections on personal computers. It will not yet be intended as a music player, but rather as a library from which users can select music and send it to their favorite player.

APPROACH

For every piece of music in the program, information is fetched from internet. For every artist there's a list of similar artists and a list of genres. For every song some audio features are collected.
The artists, albums and songs are represented as nodes in a network. The connections between the nodes are based on the genres of the music and the similarities between artists. A force based graph determines the positions of the nodes on screen.

Data source

The information about the music comes from a combination of different sources. When an MP3 file is imported into the program, the ID3 tags for *artist name*, *album title*, *year* and *song title* are read. (The genre tag is not used, because it's too unreliable due to the fact that people often change it to fit into their personal classification system.)
Besides this basic information, the most important attributes that people want to retrieve and browse their music by are *musical similarity*, genre, tempo and *mood* [Vignoli 2004]. *Mood* is not included in the system for reasons mentioned earlier. The other attributes are fetched from the internet using the information from the ID3 tags. Two different online databases are used:

- **Last.fm** collects data from Last.fm users. Its main functionality is a recommender system based on what the users listen to on their personal computers, portable music players and internet radio stations. Besides that, the database contains user-generated tags for every artists and some basic information about the music.

- **The Echo Nest** crawls texts on websites like blog posts, album reviews and tracklists of mixtapes. It counts how many times artists are mentioned in the same text, and which terms are used to describe these artists. Besides that, The Echo Nest has a database of about 30 million songs that have been analyzed to extract musical features like tempo.

These two databases were chosen because they are the most extensive in their sort. Last.fm is a user-based source and The Echo Nest is an expert-based source, so combined they cover a lot of information. They contain both absolute properties of the music (tempo, genre) and relative properties of the music (similarities to other artists/songs).

The figure below shows which information is gathered about the music.

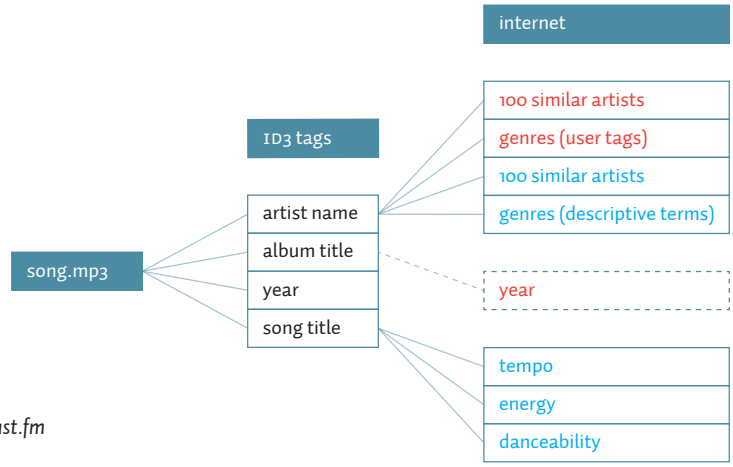

RED
data from Last.fm

BLUE
data from The Echo Nest

The *energy* and *danceability* of a song are calculated from a range of audio features. These two attributes are included as a replacement for the *mood* attribute.
If the ID3 tag of a song doesn't include a value for *year*, it is looked up in the Last.fm database.

Last.fm offers the functionality to find similar songs for a song, but in practice that only works for the most popular songs, so it's not suitable for an average collection. The Amazon API offers the functionality to find 10 similar albums for an album. The chance that any of those 10 albums occur in the same music collection is very small, so unfortunately that function is not useful either.

Representation

It was chosen to use an iconic representation to show the music. Icons provide clear relationships and a better understanding of what they actually represent.

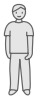

Every artist in the music collection is represented by a little man on the screen. The size of the man represents the combined duration of the songs it contains: the longer the duration of the songs by the artist, the bigger the man is.

Every album is represented by a 30'' vinyl record. The size represents the combined duration of the songs it contains, and the brightness of the color represents the year in which the album was produced: the older the album is, the greyer the color.

Every song is represented by a 7'' vinyl single. The brightness of the color represents the tempo, danceability and energy of the song. The higher the tempo/danceability/energy, the brighter the color. The size of the icon represents the duration of the song.

Every genre is represented by a label. The size of the label represents how many artists belong to this genre.

Projection technique

To translate the musical properties into positions of the music icons on sceen, a combination of two projection techniques is used:

› a **graph structure** to show the relations between similar music
› a **query map** to show to which genre the music belongs to

The artist icons are connected to any similar artists thatare in the collection, and to the genres they belong to. This means that the position of every artist in the collection is determined by both its relative properties and absolute properties.

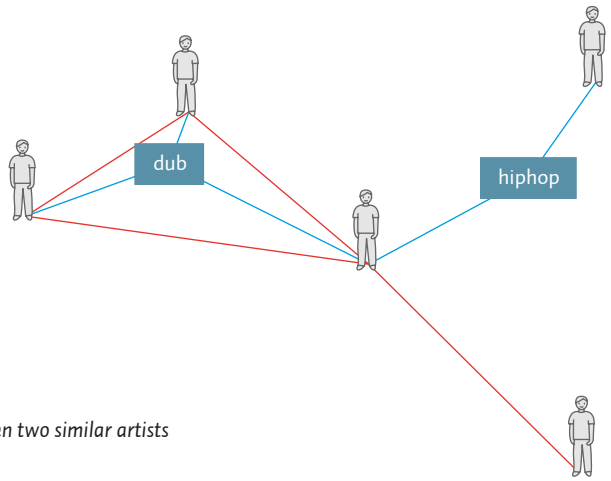

RED

connection between two similar artists

BLUE

connection between an artist and a genre

To avoid visual cluttering, only the 20 most prominent genres are displayed. A genre label is only displayed if there are at least two artists in the collection that belong to that genre with a certain probability.

Album and song level view

In addition to the artist level layout as described above, the program contains two other levels in which the user can view the collection: on an album level and a song level view. They're built up in the same way as the artist level, but have some extra connections. In the album level view, there are connections between albums by the same artist. And in the song level view, there are connections between songs from the same album, and songs from the same artist.

Layout

To determine the layout of the network on screen a force based graph is made, in which the music icons and the genre labels are nodes, and their connections are springs. The nodes repel each other, but are kept together to a certain extent by the springs. The more similar two artists are, the more stongly they are attracted to each other. And the more an artist belongs to a certain genre, the more it is attracted to that genre label.
The force based graph, combined with a smart importing method, makes sure that the layout of the nodes is optimal (i.e. as few crossing edges as possible and proportional edge lengths).

Interface features

The program has a zoom function. If you zoom in while you're in album level view, the songs on the albums are revealed. If you zoom in while in artist level view, the albums by the artists are revealed. When you zoom in even further, the songs on the albums are revealed as well.

The program has a filter function that allows the user to search for specific music. Filters can be added on top of each other, for example if you want to search for song titles that contain the word "version" and are produced before 1990.

Playlists can be made by manually selecting an artist/album/song, or by taking all the search results at once and adding them to the list. Exporting the playlist generates a M3U file that can be played in any music player.

The system includes a preview function that can play a 10-second fragment of a song. This function is merely intended for testing purposes: it enables the user to find out whether a song is what he or she expected.

USER TESTS

To find out how effective the system is, it was tested on eight people. First they were asked to perform a list of tasks within the program, to force them to learn about its functionalities. After this session the users were given a week time to try the program on their own computers at home. This enabled them to discover the system more freely and exprerience how it works on the longer run. After that week, the users were asked to fill in a questionnaire about their experience with the system.

During the first test session, the test users were asked to load some of their music into the program, and then perform a list of tasks. While they were doing the tasks it was observed how they used the software. A summary of the observations:

- In the first round of tasks, when people were asked to find specific artists/albums/songs, most users used the search function.
- In the second round of tasks, when people were asked to find certain kinds of artists, seven out of eight users used the artist-level visualization to find these artists by sight.
- In the third round of tasks, when people were asked to find a number of songs to be played at specific occasions, people searched in different ways. Four out of eight used the song node colors to find the music they'd probably want to hear. Two of the users first thought of a specific artist/album/song and then used the search function to find it. Two of them just browsed the visualization on screen.
- Three people used the play function to check whether the found songs were good, others just assumed that the songs were allright.
- Seven out of eight users didn't understand the filter function without explanation beforehand, but with explanation they all did.

An list of the tasks the users had to do is included in appendix A.

After trying the program at home for a week, the users were asked to fill in a questionnaire. A summary of the answers:

- Three out of six users said they got a better overview of the scope of their collection, mostly because of the genre labels.
- Six out of six users said they discovered new relations within their collection.
- Six out of seven users said they were more likely to discover music they'd forgotten they had. Three users mentioned that when they browse through their list of artists, they usually skip some artists because they don't remember what kind of music it was. Instead of taking the effort of finding out what music is it, they go for a quick safe option: the music they already know well. Because the new program offers more clues on what kind of music an artist is, the users said they would listen to those 'unknown' artists more often.
- In general, all of the users thought the position of the music icons made sense. Three people expressed that they didn't always agree with the connections between genre labels and artists.
- Five out of seven people used the artist level the most. The most important reason was that they're used to retrieving music by artist.
- In general, the users found it easy to find music in a specific genre and music for a specific occasion.
- Six out of seven people said that the song icon colors for tempo, *energy* and *danceability* often don't match.

An extensive report of the questionnaire results is included in appendix B.

EVALUATION

The questions that were formulated to evaluate the system can now be answered based on the user tests:

Can people find the exact artist/album/songs they are looking for?
Yes, the user tests show that in general it is easy to find a specific piece of music.

Can people find the kind of songs they would like to hear, for example songs that fit their mood?
Yes, the user tests showed that the program makes it easy for the users to find the kind of music they'd like to hear, in different ways. Some test users said that it enabled them to easily find music that belongs to a specific genre or mix of genres. Some test users said that it enabled them to easily find music they'd like to hear on a specific occasion.

Do people get a better overview of the scope of their music collection?
Half of the test users said the program gave them a better or different overview of their collection. All of the test users claimed to have discovered new relations within their own collection.

Are people more likely to discover music in their collection they hadn't heard before?
Yes, the user tests show that most of the test users are more likely to discover music in their collection they hadn't heard before. Even more test users said they discovered music that they'd forgotten they owned.

Generally speaking, the user tests showed that the new system helps the users to find the music they want to hear. The visualization provides a truly new view on the information that can't be easily conveyed in text. It shows the positioning of a piece of music with respect to other pieces of music, and at the same time it shows the positioning of the music with respect to genres. It provides a good overview of the music at a glance and offers new insights into the collection.

FUTURE WORK

There are still some features to be included in the system.

More musical properties could be included. Useful properties of songs could be *play count*, *popularity* (for example measured by number of *listeners* on Last.fm) and the time of day on which a song is generally played.

To improve the visualization, a heatmap-like image could be displayed in the background. It could provide a better visual clue about the density of connections between nodes in the network.

The user interface design needs improvement in order to become suited for everyday use. Not much attention was paid to the interface design because it was not in the scope of the research question. The program could be extended to be a music player in order to remove the gap between the visualization and user's regular music player.

REFERENCES

› Dittenbach, M., Neumayer, R. and Rauber, A. *PlaySOM: An alternative approach to track selection and playlist generation in large music collections*. AVIVDiLib 2005.

› Dalhuijsen, L. and van Velthoven, L. *MusicalNodes - the digital music library*. Electronic Visualisation and the Arts 2010.

› Hitchner, S., Murdoch, J. and Tzanetakis, G. *Music browsing using a tabletop display*. ISMIR 2007.

› Julia, C. and Jorda, S. *Songexplorer: a tabletop application for exploring large collections of songs*. ISMIR 2009.

› Knees, P., Pohle, T., Schedl, M. and Widmer, G. *Exploring Music Collections in Virtual Landscapes*. IEEE Multimedia 2007.

› Lillie, A. *MusicBox: Navigating the space of your music*. Master's thesis Massachussets Institute of Technology 2008.

› Rauber, A., Pampalk, E. and Merkl, D. *Using psycho-acoustic models and self-organizing maps to create a hierarchical structuring of music by musical styles*. ISMIR 2002.

› Vignoli, F. *Digital music interaction concepts: A user study*. ISMIR 2004.

› Vignoli, F., van Gulik, R., and van de Wetering, H. *Mapping music in the palm of your hand, explore and discover your collection*. ISMIR 2004.

Appendix A: Test session tasks

These tasks where created automatically for every test user, based on the music they loaded into the program. The blue words in the text below are placeholders for collection-specific items. For example, if a user had the artist Ghostface Killah in his collection, the first task could be '*all songs by Ghostface Killah*'.

Create a playlist with the following music:
- all songs by artist
- all songs from the album album by artist
- the song song by artist
- the song song
- all songs from an album from year
- all songs from an album from after year

Create a playlist with the following music:
- all songs by an artist thats unlike any other artist in the collection
- 5 songs you would describe as genre
- 5 songs you would describe as a mixture of genre and genre

Create a playlist with 5 songs you would like to hear on sunday morning
Create a playlist with 5 songs you would like to hear before going out on saturday night
Create a playlist with 5 songs you would like to hear while at work/study

Appendix B: Questionnaire

Seven users returned the questionnaire. One of them answered only half of the questions due to circumstances (that's why some questions are answered by six people, and some by seven).

A list of the questions and a summary of the answers people gave:

Did you get a better overview of the scope of your own music collection than before?
Three out of six users said they get a better overview of the scope of their collection, mostly because of the genre labels. One of the test users called it a 'rather different kind of overview'

Did you find new relations within your collection?
Six out of six users answered yes. Some of them said it was because of the genre labels: they didn't know what kind of genre a certain piece of music was, and they do now. Some of them said it was because of the connections between similar artists that they hadn't considered before.

Do you think you're more likely to discover music in your collection you hadn't heard before?
Four out of seven users answered yes. Most of them know what music is in their collection. However, they do forget they have it (see next question).

Do you think you're more likely to discover music you forgot you had?
Six out of seven users answered yes. Three users mentioned that when they browse through their list of artists, they usually skip some artists because they don't remember what kind of music it was. Instead of finding out what music is it, they go for a quick safe option: the music they listen to often. Because the new program offers more clues on what kind of music an artist is, the users said they would listen to those 'unknown' artists more often.

In the artist level view, do you think the positions of the artists in relation to the genre tags make sense?
In the artist level view, do you think the positions of the artists in relation to other artists make sense?
In the album level view, do you think the positions of the albums in relation to the genre tags make sense?
In the album level view, do you think the positions of the albums in relation to other albums make sense?
In the song level view, do you think the positions of the songs in relation to the genre tags make sense?
In the song level view, do you think the positions of the songs in relation to other songs make sense?
In general, all of the users thought the position of the music icons made sense. Three people expressed that they didn't always agree with the connections between genre labels and artists. For example, in one user's collection Bob Dylan was connected to alternative rock. In Bob Dylan's term list, this is the 6th most used descriptive term, but indeed not the first thing you would think of. However, the icon for Bob Dylan was connected to this genre because apparently there where at least 2 other artists in the collection who's music was also described as alternative rock. If there would be more music in the collection, Bob Dylan would probably be connected to other genres as well (like classic rock, folk or singer-songwriter), and the positioning of his music would become clearer.
Another user suggested using more global terms when there are only few artists in the collection, so for example electronic instead of breakcore, but the problem is that the system cannot (yet) tell which term is more global and which one is more specific.

Which view level did you use most? (artist, album or song or any subset of these three levels) Why?
Five out of seven people used the artist level the most. The most important reason was that they're used to finding music by artist. Another reason that was mentioned was that the artist view is the most clear.

How difficult was it to find a specific artist?
How difficult was it to find a specific album?
How difficult was it to find a specific song?
All users said it was easy to find a specific artist/album/song. Only two people said they found it difficult to find a song. That was because if the collection was big, and/or they were zoomed in, they couldn't find the highlighted song.

How difficult was it to find music in a specific genre?
How difficult was it to find the kind of music you would like to hear at a specific occasion?
I got varying answers to these questions. Three out of six people said the first was easy and the second was neutral. Two out of six user said the other way round: the first one was neutral, the second easy. One person said both were super easy.

Do you think the color representing song tempo makes sense?
Do you think the color representing song energy makes sense?
Do you think the color representing song danceability makes sense?
Do you think a combination of any of the properties above makes sense?
Most people answered that the sometimes thought the values for these properties matched, but often not. This is what I also exprienced in my own collection, but I was wondering if these values would perhaps make sense for other kinds of music. Apparently not.

Do you have other comments?
Some people mentioned missing a play function. One person mentioned wanting an intelligent playlist function that automatically adds certain kinds of music. This was also mentioned by another user during the first test sessions. One person mentioned wanting to add the play count of a song as property.