

CloudMonster: Support Flexible Browsing and Searching within Music Collections

Ya-Xi Chen, Matthias Hoyer, Andreas Butz

Media Informatics, University of Munich
Amalienstr. 17, 80333 Munich, Germany

yaxi.chen@ifi.lmu.de, matthias.hoyer@campus.lmu.de, andreas.butz@ifi.lmu.de

Abstract. Studies in the field of Music Information Retrieval (MIR) have shown that users would like to use multiple criteria simultaneously and additional criteria besides the widely used metadata of artist, album and song. We present our prototype named CloudMonster, which supports flexible browsing and searching within music collections using multiple criteria. In a preliminary evaluation the added flexibility for browsing and searching was highly appreciated and we received valuable suggestions for future research.

1 Motivation and formative study

For Music Information Retrieval (MIR), there are additional criteria besides the widely used artist, album and song. These criteria, for example, similarity, mood and social activities, are less common but nonetheless desired [1]. Cunningham et al. [2] suggested that functionalities beyond explicit searching should be offered, to allow users to find unexpected but acceptable results. Some researchers have already built visualizations based on non-standard criteria, such as similarity [3] and mood [4].

In order to learn more about the actual criteria people wish to use, we conducted a formative study. We recruited 36 college students, 25 male and 11 female, with an average age of 22 years and investigated their current browsing behavior. From this, we learned that the most desired criteria for organization seem to be genre, year and cover art, while for playlist generation, they are gender, tempo, instrument and listening history. By combining the results of our survey with existing research, we selected the set of most desired criteria and integrated them into a prototype named CloudMonster for browsing and searching within music collections.

2 CloudMonster

CloudMonster is implemented based on Prefuse (<http://www.prefuse.org/>). The basic metadata, such as genre, duration, released time and cover art are extracted from ID3Tags. The content-based similarity is derived from low level features extracted with jAudio (<http://jmir.sourceforge.net/jAudio.html>). Contextual information such as artist popularity and similarity are derived from Last.fm (<http://www.last.fm>).

Figure 1 shows the CloudMonster interface, in which songs are represented as either nodes or cover art. Upon mouse-hover, more song detail is shown (see figure 3(b)). The whole collection is displayed with genre-, popularity- and similarity-based *initial views*. Five *color coding schemes* are provided: for ordinal parameters, each color presents one category such as artist, genre and gender and for quantitative parameters such as popularity and year, the transparency represents the popularity or recency.




Fig. 1. CloudMonster with genre-based view and genre color coding. (A) Overview. (B) Keyword-based searching. (C) Saved playlist. (D) Saved criteria. (E) Genre histogram. (F) Cover art. (G) Initial views. (H) Reset button. (I) Color scheme. (J) Criteria list. (K) Song list.




Fig. 2. (a) Popularity-based view with popularity color coding. (b) Similarity-based view with artist color coding.

Based on the Dust&Magnet metaphor[5], each criterion works as a *magnet* and attracts matching songs while repelling differing ones. The user can create or delete a magnet by double clicking a criterion in the criteria list. The magnet value can be set in the criteria list and its strength is presented as its size, which can be adjusted by mouse wheel. The user can then adjust the magnet position freely and the songs will move accordingly. For example, three magnets were chosen in figure 3(a): BPM, gender and genre. The values were set as 136-160, female and R&B. These magnets were placed in a triangle. The matching songs were attracted in the center and shaped in a circle. The current configuration and visualization can be saved for future reuse.

Similarly, one *song* can serve as an *example* and attract similar songs (see figure 3(b)). The weights for content- and context- similarity can be adjusted by a weight slider.

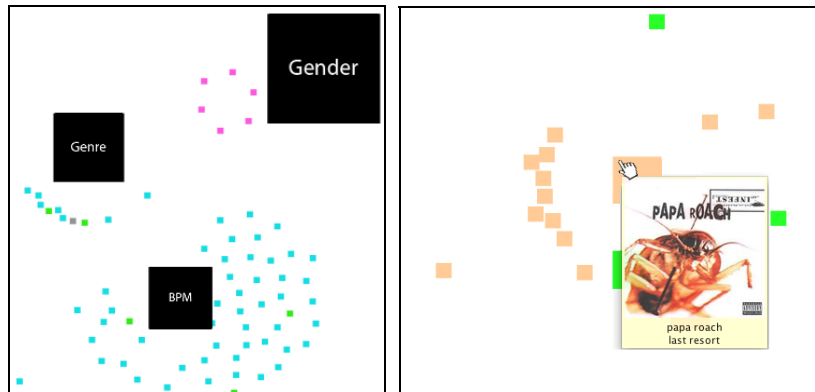


Fig. 3. (a) Multiple-magnets-based search. (b) Example-based search.

The user can *generate a playlist* by simply drawing a rectangle in the graph view. All the songs inside will then be included in a playlist, which can be saved. Once a saved playlist is chosen, the graphic for the moment this playlist was created will be re-visualized, which also improves the system reproducibility and speeds up learning.

3 Preliminary evaluation

We recruited 12 users at (delete for blind review), 7 female and 5 male with an average age of 24 years, who regularly use music player software. The evaluation was conducted with the participants' own music collections. After a brief tutorial and freely playing around, participants were asked to conduct 5 tasks concerning browsing, searching and playlist generation, and then filled out a post-questionnaire.

The tasks were: Locate some personal favorite songs; from these, filter out faster songs; search for the songs similar to an example song; search for songs of the favorite genre and released in the same year; generate a playlist for a party.

Post-questionnaire: After completing the tasks, the participants were asked to score the initial views and color coding schemes (1 for not useful at all and 10 for very useful). The average scores for genre-, popularity- and similarity-based views were 8.75, 8.23 and 7.75. For the color coding schemes, the average scores for genre, popularity, gender, artist and year were 8.25, 8.16, 8, 7.8 and 7.125. The average score for the overall impression was 8.0. For helpfulness, enjoyments, feeling of control and understandability the average scores were 8.75, 8.5, 8, 8, and 7.5.

Results and implications: In task 1, all the participants conducted a keyword-based search. In task 2, they firstly conducted keyword-based search for the artist, and then used a BPM-magnet to attract songs with higher BPM value. The usage of keyword-based search illustrated its advantages with large collections. In task 3, the participants firstly chose the similarity-based view, and then chose one example song

from the song list. This song was highlighted as an example in the graph and similar songs were attracted automatically. In task 4, the participants used the genre- and year-magnets. Without being prompted, all of them were able to make the right judgment of matching songs, which implied the understandability of the magnet metaphor. In task 5, the most frequently used criteria were BPM, genre, popularity and instrument. Generally positive feedback was received in the post-questionnaire. Although all the participants claimed that the magnet metaphor was easy to understand, the average score for the overall understandability was lower because of the attribute inconsistency in the views, color coding schemes and criteria list.

4 Conclusion and future work

In this paper we present CloudMonster, which supports flexible browsing and searching within music collections. CloudMonster received positive feedback in an initial evaluation and its multiple search options were highly appreciated. CloudMonster can be easily extended for example for browsing in large online music dataset or generalized for broad multi-criteria search. In our future work, we will address the problems revealed in the evaluation, for example, improve the systems consistency, allow duplicate magnets and multiple example songs, and enhance the histogram for color coding. Then we will conduct a formal user study concerning systems scalability, data reliability and look in detail at the behavior differences between different visualizations, such as list and graph.

5 Acknowledgement

This research was funded by the China Scholarship Council (CSC) and by the German state of Bavaria. We would like to thank Eric Rademacher and Ben Blaha for their dedication in the initial design and thank the participants of our user study.

References

1. Vignoli F. Digital Music Interaction concepts: a user study. Proc. of the 5th International Conference on Music Information Retrieval (Barcelona, Spain, 2004). ISMIR '04.
2. Cunningham S, Reeves N. and Britland M. An Ethnographic Study of Music Information Seeking: Implications for the Design of a Music Digital Library. Proc. of ACM/IEEE joint conference on Digital libraries (Houston, TX, USA, 2003). JCDL '03.
3. Pampalk, E. Islands of Music: Analysis, Organization, and Visualization of Music Archives. Journal of the Austrian Soc. for Artificial Intelligence, 2003.
4. Gulik R van, Vignoli F. Visual Playlist Generation on the Artist Map. Proc. of the 6th International Conference on Music Information Retrieval (London, UK, 2002). ISMIR '05.
5. Yi J. S., Melton R., Stasko J., Jacko J. Dust & Magnet: multivariate information visualization using a magnet metaphor. Proc. of the international conference on Information Visualization (London, UK, 2005). IV '05.