

MusicLand: Exploratory Browsing in Music Space

Heidi Lam⁺

Tamara Munzner⁺

University of British Columbia

ABSTRACT

Most existing search tools based on query terms focus on direct search activities, where users are assumed to have a clear and precise idea about their search targets. However, if users are uncertain about their targets, they will need to define or refine them with a succession of multiple related queries. Current search tools do not adequately support this process. MusicLand is designed for exploratory visual browsing of digitized music collections annotated with metadata. It supports the gradual refinement of queries by showing visually structured search results connected with a query history trail, and providing suggestions for new query terms. Semantic zooming maximizes information density, and animated transitions between display states help the users maintain their mental map of the relationship between new and old queries.

1 INTRODUCTION

The information seeking behaviour of users can be roughly classified according to their certainty of the target, for instance by knowing the name or attributes of a file, and their certainty of its location, for instance its full path within a file system. When users are certain of the target location, they can navigate or “jump” directly to access the target. If the user does not know the target location, a search tool is required. However, when the user knows neither the location nor the precise target, the information seeking behaviour is different. It can be characterized as exploratory, to discover a local neighborhood of interest; and browsing, to surf interesting neighborhoods in detail.

Since most search tools focus on single, precisely targeted queries, exploratory browsing is not adequately supported. While users are still required to input query terms to narrow their information space to the regions of interest, browsing-oriented visualization systems should support gradual refinement of queries. The requirements for an ideal browsing visualization solution are:

1. Provide context to allow users to interpret the query results based on their input terms and address questions like “where am I?” and “what am I looking at?”
2. Guide navigation to help users refine existing queries into new queries: “given where I have been, where should I go next?”
3. Assist refinement or defining of target based on available choices: “what *can* I look for?”.

This paper introduces a new system called MusicLand as an investigation of how computer interfaces can support exploratory browsing in the domain of annotated digitized music. Figure 1 shows the MusicLand interface.

2 MUSICLAND

MusicLand explicitly supports exploratory browsing with an animated visual representation of the history of the query refinement process, and by suggesting new search terms based on

the annotations in the music. The visual display can also be used to launch a music player so that the user can hear the music.

2.1 Query result display

The structure of the visual representation of a query is based on a Venn diagram, as shown in Figure 1. MusicLand supports queries of up to three terms, a limit based on a study showing that 97% of all queries use three or less terms [1]. Query results are slotted into one of the seven regions on the 3-term Venn diagram. To maximize information density, MusicLand uses semantic zooming with four zoom levels, illustrated with Figure 1d: (1) full, in the central black region; (2) partial, only the artist name and album title, in the lower tan region; (3) minimal, only artist name, in the turquoise and light green regions; and (4) nil, only the number of matching results, in the cyan, purple, and pink regions.

The display uses colours to further reinforce the spatial relationship between the regions. The three individual query terms are coded by three colors: desaturated red, green and blue. The intersection regions for results that match multiple query terms have perceptually additive colours; for example, Figure 1d shows tan overlap between the pink *Electronica* and light green *Melancholy* regions.

2.2 Query formulation

MusicLand supports easy query refinement with a popup menu containing a list of selectable descriptors for the selected album. Selecting of an addition one to two descriptors will initiate another query. In MusicLand, these descriptors include the artist, genre, style and tone of the music*.

2.3 Query trail

To better support query formation and navigation, MusicLand retains a visual history of all previous queries. Since every new query originates from old queries, this relationship is explicitly depicted by a link between the two, using the colour of the region from which the query originated as shown in Figures 1a and 2d. This trail provides context to address the question “given where I have been, where should I go next?”.

2.4 Animation

Animation provides transitions between different semantic zoom levels and conveys the idea that new queries are extensions of their originating queries. Clicking on any region will increase its zoom level to the highest displayable level. When zooming, the space reallocation algorithm is similar to TableLens [2], where zooming into a cell in a table reduces the sizes of its neighboring rows and columns. Priority is given to regions satisfying larger number of query terms. New query animation links the old and the new query by showing how the new query “grows” out from the old.

⁺ {hllam, tmm}@cs.ubc.ca

* The English tags describing the tone of the music were generated with the MILQ machine-learning system. (<http://www.cs.ubc.ca/~ebrochu/milq.htm>)

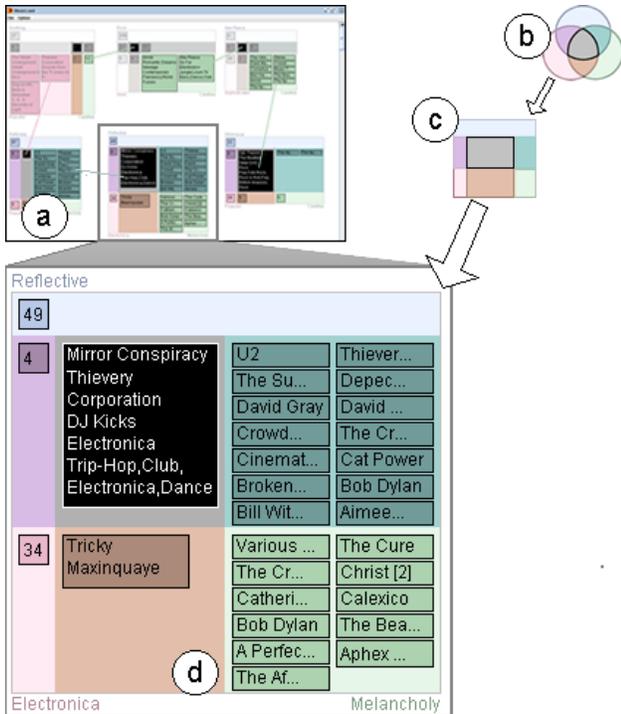


Figure 1. The MusicLand Interface. (a) overall view with multiple queries; (b) the original, and (c) the rectangularized 3-term Venn diagram; (d) single query display.

3 SCENARIO OF USE

Imagine a person browsing through a large annotated music collection at a music store, with per-album tags that include descriptive query terms like *Soothing*, *Peaceful* and *Carefree*. Figure 2a shows the result of that first query. Figure 2b shows the zoomed-in display when the user explores the search results by clicking on the pink *Peaceful* region. Seeing his favorite band Thievery Corporation in the *Peaceful* region, he plays the album by double-clicking in its box. He likes what he hears, and wants to find other music with a similar mood. He uses the popup menu in that box to reveal all the tags for that album as shown in Figure 2c. These tags are potential query terms. He picks the terms *Melancholy* and *Reflective* to initialize another query, with the result shown in Figure 2d. He notices the new query is visually linked to the old with a colored link by the *Thievery Corporation* album he originally selected. Delighted to find another album with the moods *Peaceful*, *Melancholy*, and *Reflective* amongst the retrieved results in the new query, he selects and plays the *Belle & Sebastian* album.

4 FUTURE WORK

MusicLand is a first prototype for explicitly supporting exploratory browsing in annotated music collections. The current interface has limited scalability: navigating through long browsing sessions would require a large amount of scrolling. Also, the three query colours are reused in new queries. While such an encoding scheme is appropriate for shared search criteria among the queries, it can result in a misleading perceptual grouping of objects when the query terms are changed. MusicLand partially addresses this issue by fading and shrinking unused regions, as shown in the early

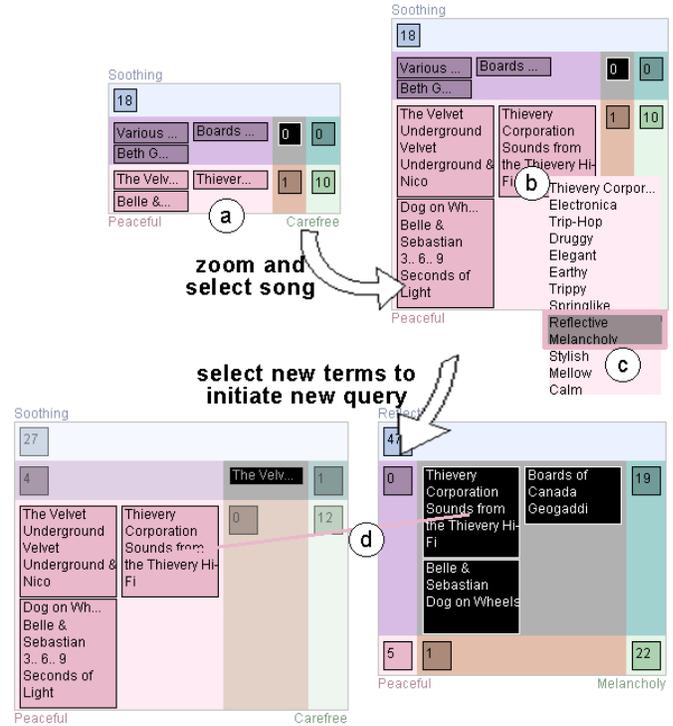


Figure 2. Scenario of use. (a) initial query; (b) zoomed in view of query; (c) display of selectable attributes of the music record; (d) display of new query

queries in Figure 1a. Our future plans are to address visualization issues of the current MusicLand, to further study the requirements of browsing versus searching behaviour, and to investigate potential roles of visualization interfaces in domains such as e-commerce and multimedia digital libraries.

REFERENCES

1. M. Hertzum and E. Frøkjær. Browsing and querying in online documentation: a study of user interfaces and the interaction process. *ACM Transactions on Computer-Human Interaction* 3(2), 131-161, 1996.
2. R. Rao and S. K. Card. The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus + Context Visualization for Tabular Information, *Proceedings of SIGCHI '94*, 318-322, 1994.