MusicLand: Exploratory Browsing in Music Space

Heidi Lam[+]

University of British Columbia

ABSTRACT

Most of the existing query-term based search tools focus on direct search activities, where the user is assumed to have a clear and precise idea about his targets. However, if the user is uncertain about his targets, he will need to define/refine them with multiple related queries. Such query behaviour, or exploratory browsing, is not well supported by current search tools. MusicLand is designed for exploratory browsing using music album data. Users input query terms to select a potentially interesting subset of the database. To provide context, the retrieved results are visually clustered based on the query terms, and are displayed on a rectangularized Venn diagram. The Boolean relationships between these regions are visually reinforced using colour encoding and perceptual layering. Refinement of old query is guided by a history of the user's query evolution, along with a list of potential query terms based on the user's current music album of interest. Semantic zooming maximizes MusicLand's display capacity. Animation helps link the zoom views, and the new and old queries. A formative evaluation revealed encouraging responses, suggesting the validity of MusicLand's design principles.

CR Categories: I.6.9.c Information visualization, H.5.2 User Interfaces, H.5.2.f Graphical user interfaces, H.2.8.h Interactive data exploration and discovery

Keywords—information visualization, exploratory browsing, semantic zooming, focus+context, music database query

1 INTRODUCTION

Information space, may it be personal or shared, has increased both in size and complicity thanks to the advances and popularization of computer technology. This double-edged sword grants an average computer user access to an enormous amount of information, and ironically, also impedes the process of information retrieval. Previous efforts in solving the problem of information retrieval in large and complex information space using information visualization techniques have largely focused on displaying query-based results of text documents where users are assumed to have a relatively clear and precise idea about the piece of information they wish to find with the tool. While that may be true in most office situations, there are occasions where users' targets may not be as well defined.

Users' information seeking behaviour can be roughly classified according to their certainty of the target (e.g., file names or file attributes), and of its location (e.g., file path). When users are certain of the target location, they can navigate or "jump" directly to it to access the target. Using a scenario of e-commerce where the user wishes to find a particular music album, he can navigate from the store's home page to the music subsection and then, to the album's web page. Alternatively, he may have visited and bookmarked the page of the album, and can directly access it via the predefined bookmark. On the other hand, if the target location is unknown to the user, a search tool is required to retrieve the location of the target. Continuing with the e-commerce scenario, the user will need to input keywords that describe the record (e.g., the name of the album) to a web search engine to retrieve the URL of the album page. However, when neither the location nor the target is known to the user, the information seeking behaviour is different, and can be characterized as exploratory (to discover a local neighbourhood of interest), and browsing (to explore interesting neigbourhoods in detail). In the e-commerce scenario, the user may be looking for a gift for a friend. While he may have an idea about his friend's musical taste (e.g., Classical music), he may not know exactly which album would be the best choice. In that case, he may first navigate to the Classical section of the online store, locates the area of interest based on available selections, and browses until he finds his target.

As seen from this simple scenario, the goals and task of exploratory browsing differ from those of searches, where the user's target is known and static, and his goal is to find the information in question as quickly as possible. On the other hand, with exploratory browsing, the user's target is relatively imprecise, and will likely change in the course of the query. In the e-commerce scenario, the user may come across an album that plays J. S. Bach's music in the Jazz style, and decides to get his friend a Jazz album instead. In addition, the user will most likely put more

[+] hllam@cs.ubc.ca

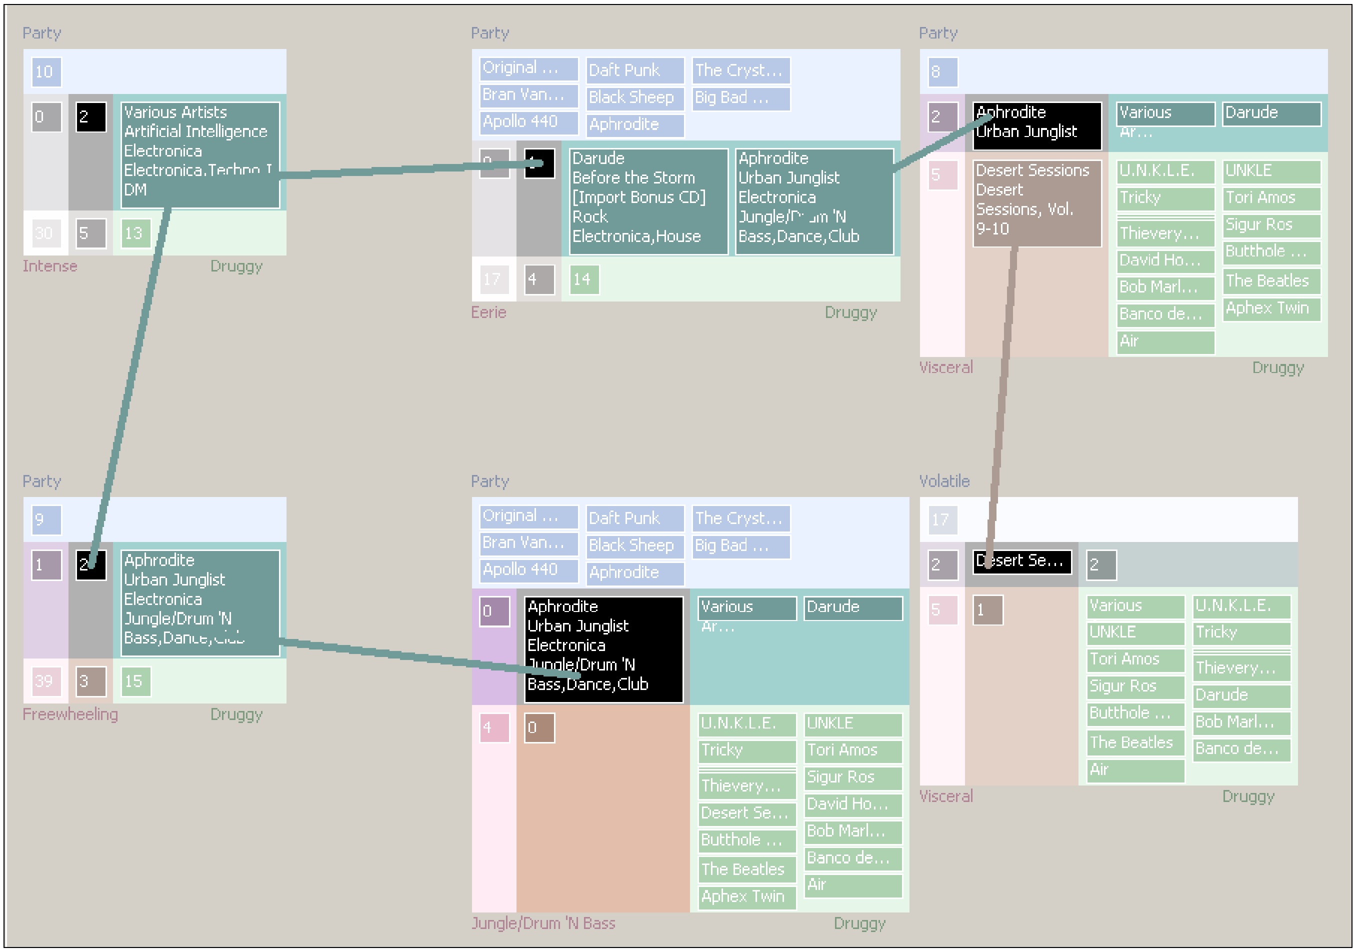

Figure 1. MusicLand overall view with multiple queries.

emphasis on the query process than in the case of direct, simple searches (i.e., "getting there is half of the fun").

Since most search tools focuses on single, precisely targeted queries, exploratory browsing is not adequately supported. While the user should still be required to input query terms to narrow their information space of interest, the visualization should better support refinements of the user's current and past queries into future queries. Basically, the ideal visualization solution should,

1. Provide context to allow users to interpret the query results based on their input terms (i.e., "where am I?" and "what am I looking at?".
2. Guide navigation to help users refine their current and past queries into new queries (i.e., "given where I have been, where should I go next?")
3. Assist refinement of target based on availably choices, since user's target may change in the course of the quest (i.e., "what *can* I look for?").

This paper introduces a new system called MusicLand (Figure 1) as an attempt to study how computer interfaces can support exploratory browsing in the domain of music album data.

The paper is arranged as follows. In Section 2, I will survey related work in the areas of document space visualization. Section 3 details the design and implementation of MusicLand, followed by a scenario walk-through in Section 4. I will then discuss the strength and weaknesses of MusicLand in Section 5, along with the lessons learned from working on the MusicLand project. The paper will end with a list of future work in Section 6 and a conclusion.

2 RELATED WORK

There are two main approaches to visualizing document space: visualizing the whole space, or visualizing a small subset of the space based on user-input query terms. For both approaches, the visualization may adopt a temporal or spatial metaphor. Since MusicLand is a query-based tool, I will only briefly outline a few related work that visualize the entire information space.

The main purpose in visualizing the whole document space is to obtain an overview of the space to discover the main themes of the documents. Based on Kohonen's feature map, Lin mapped a high dimension document space to 2D space to visualize documents on the user's computer [15]. Galaxy uses the starfield metaphor to display text documents in 2D [21], and Themescape [21] and BEAD [4] use the landscape metaphor in 3D space. While these visualizations use spatial metaphors to show snapshots of document themes, ThemeRiver displays them in a time-line using the river metaphor [11].

For keyword-based queries, the most popular type of display is text summary, where a small sample of the document containing the query terms is displayed along with the document's location (examples includes most web search engine interfaces and Stuff I've Seen [7]). Despite capturing a large amount of information with the text snippets, this approach requires reading, a potentially mentally intensive process, and does not provide layout information of the document to take advantage of our powerful visual processing abilities. In light of this, Woodruff et al. proposed Enhanced Thumbnails that show query terms as readable callouts on a thumbnail of the document [22]. To better delineate the relationships between the retrieved documents, Lighthouse clusters retrieved documents from the search engine and display them in 2- or 3-D space [14], and Sparkler shows the ranking of the documents on concentric circles to allow visual comparisons of multiple query results [10]. In the music domain, MusicPlasma™ (www.musicplasma.com) retrieves the names of the artists with similar music as the originally user entered artist, and displays them as orbits around the original entry in 3D space.

A number of systems visually group the documents based on the query terms in the spirit of the Venn diagram. VIEWER [3] lists all Boolean relationships between query terms and the number of documents retrieved that satisfies these conditions. User can select any of these conditions and display a linear list of documents summaries. Cougar [12] displays the number of documents retrieved for each region on a 3-term Venn diagram on the diagram. InfoCrystal [20] abstracts the Venn diagram into an iconic display, and can handle more than three query terms. VISER [23] places the query terms on the parameter of a circle, and arranges thumbnails of the retrieved images inside the circle according to their similarity to these query terms.

Instead of arranging the retrieved results spatially, an alterative approach is to display them along a time line. An example is Milestones in Time [18], which provides personally significant events as temporal landmarks on a time line as contexts for users to assess the relevance of the retrieved documents. Since both the spatial and temporal views may be of use, there are systems that provide both. For example, InfoSpace [17] provides coordinated spatial (based on document similarity) and time-line (based on time of the document's last retrieval) views of the query results.

While these visualization systems are designed for single static queries, there are systems that support dynamic queries. In dynamic queries, the values of existing query terms can be modified interactively, and the systems provide immediate feedback. Examples of such systems include HomeFinder and FilmFinder [1].

3 MusicLand

MusicLand is a system that explicitly supports exploratory browsing. It does so by displaying the individual queries in the context of their query keywords in a Venn diagram like display. The display uses colours to further reinforce the relationship between the regions and to the original query terms. To better support navigation, MusicLand retains all the previous queries as a trail of the user's query history, and provides a list of possible query terms based on user's current interests to guide user's next query. This section will explain the main visualization features of MusicLand in detail.

3.1 Overall Spatial Layout

The spatial layout of MusicLand is based on the 3-term Venn diagram (Figure 2). Instead using the more familiar form of the Venn diagram (Figure 2a), MusicLand uses the rectangularized version (Figure 2c).

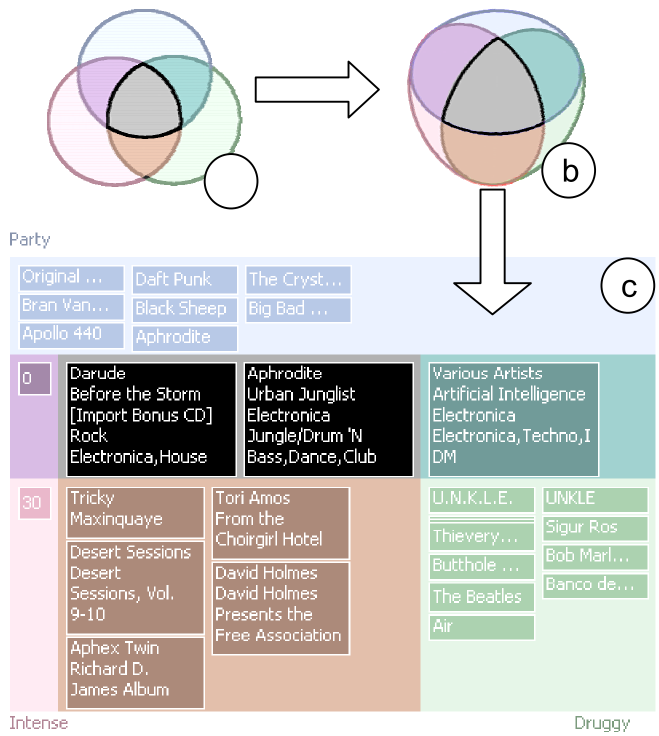

Figure 2. Individual query result in MusicLand. The spatial layout of MusicLand (c) is based on a 3-term Venn diagram (a).

The reason for choosing a graphical layout instead of the more popular list layout for MusicLand is two-fold. In the situation where the target is well defined and with an effective search engine, the user can usually find his target from the top retrieved results, and does not need to understand the overall structure of the results. In that case, even though clustering of data may help focus the user's attention to documents of interest, the cost of the more spatially expensive, and hence, lower information density graphical layout does not seem justified. However, in the case of exploratory browsing, the user does not have a precise target in mind, and the structure of the retrieved results may help him refine his target in future queries. Also, while it is relatively easy to summarize text documents with lines of text, the task is more complicated for music data. Arranging them spatially in the context of the query terms may help to characterize the less tangible music data.

For 2D displays, the 3-term Venn diagram is an attractive choice for three reasons. First, it is an effective visual display to delineate all possible Boolean relationships of the three query terms. The main purpose is not to spell out all these Boolean relationships, but to produce a simple way to cluster the retrieved results such that documents within each region are as homogeneous as possible. Second, regions in the Venn diagram are related to their immediate neighbours by at least one common query term. This feature is particularly desirable for MusicLand, since navigation from one region to another will always be perceptually "smooth". Third, while users may prefer inputting a large number of query terms, a study that looked at browsing and querying online documents showed that 97% of the free-style Boolean queries produced by their study participants could be displayed using a 3-term Venn diagram [13]. Even though the study focused on searching on text documents, it is likely that three inputs per query will be sufficient in exploratory browsing with music data, since with less specific targets, users will more likely require less instead of more input terms.

Despite these attractive features, the original non-rectangular Venn diagrams have inherently limited available display space as a container for rectangular text boxes (i.e., individual music data record). If the system attempts to maximize the total number of music data records displayed, it will be unlikely for the record boxes to line up as rows and columns, thus resulting in a display that is visually noisy and cluttered. On the other hand, the rectangularized version (Figure 2c) is a more abstracted form of the Venn diagram, and users may not be able to delineate the main regions of the diagram (e.g., where is the "Party" region in Figure 2c?). In fact, this task is already difficult when the original diagram is transformed to one with a smooth outline (Figure 2b). However, since the main purposes of the display is to cluster the retrieved results and guide navigation instead of conveying the Boolean relationships between the regions explicitly, the benefits of using a rectangular display may be a justifiable tradeoff.

3.2 Music Record Displays

To increase the display capability of MusicLand, music records are further clustered as album piles based on the nature of the music. In Classical music, compositions are frequently identified with their composers, and music listeners may have a preference for certain composers. Such personal collections may contain many pieces of compositions, but only by a handful of composers. Similar argument may be applied to other genres of music, where it is possible to cluster by artists. Other clustering parameters may be genre (e.g., Classical, Jazz) or style (e.g., Indie Rock, Garage Punk).

These album piles are depicted in Figure 3. The circled area in Figure 3a shows four albums by the same artist (Thievery Corporation) that are piled to conserve screen space. Similar technique can be used even when there is only enough space to display the name of the artist. In that case, the piled albums are depicted by horizontal lines to suggest staggered top edges of the piled records (Figure 3b).

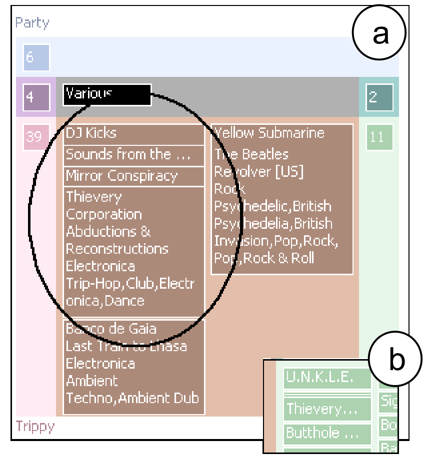

Figure 3. A single query result display showing piling of music records in the (a) full detail display, and (b) minimal display where only the artist's names are shown.

3.3 Colour

Colour encoding. The three query terms are colour-coded with the three primary colours: red, green and blue. For example, in Figure 2c, the query term "Party" is coloured blue, and the term "Druggy" is coloured green. The regions corresponding to the intersection of these terms are coloured with the perceived addition of the basic colours to indicate their relationships with the query terms and regions. For example, the region corresponding to the intersection of the terms "Party" and "Druggy" is coloured turquoise, a mixture of blue and green. The choice of using the primary colours is due to their distinctiveness.

Perceptual layering. In order to emphasize the relevance of the query regions and to reduce visual cluttering, the technique of perceptual layering is used. Perceptually, layering reduces visual clutter and organizes task relevant data elements since user's attention will be directed from the most important to the least important objects on the display due to the illusion of depth difference produced by layering. This technique has been advocated by Van Laar [24]. Briefly, all colours within the same layer must have the same or very similar level of lightness and saturation, and colours between layers must be sufficiently different (by at least 10 ΔE^*_{uv}). Also, to avoid colour interaction, the colours used in the display should be unsaturated.

In MusicLand, three main layers are created using the colour palettes created by Van Laar [24], each with a sub-layer for the music data record boxes, and another for the background:

(1) *Top most layer*: This is the region where the three query terms intersect, and is the most important region of the display. The background is a grayish colour and the music record background layer is black, the most saturated sum of the three query region colours.

(2) *Middle layer*: This consists of regions where two of the three criteria intersect. The background colours are taken from the middle palette created by Van Laar, while the music record background colours are from a more saturated palette.

(3) *Lower most layer*: This consists of regions without any intersections. The background colours are taken from the lightest palette and the background of the music record boxes are taken from the middle palette.

3.4 Semantic Zooming and Animation

Due to space constraints, it is unlikely that MusicLand can display all the music records in full details. Semantic zooming is therefore used to maximize the total information content and context of the display. There are four zoom levels in MusicLand:

(1) *Full Detail*: All the information pertaining to the music records are shown, including the album titles, the artists, the genres, and the styles;

(2) *Partial Display*: Only the artist and album titles are displayed;

(3) *Minimal Display*: Only the artists are displayed;

(4) *Nil Display*: Only the numeric value of the number of records in that region is displayed.

The zoom level of a region depends if it is in focus. In the initial query display, the most relevant region is the one where the three query terms intersect, and is in focus by default. Music records belonging to the focal region are given priority and will be displayed in full detail if spatially possible. Similarly, regions corresponding to the intersection of two query terms will be given space priority over regions corresponding to only single query terms.

In order to allow the user to explore details of the retrieved results that are not shown at the highest zoom level, user can select another region by clicking on the region, thus putting it in focus and increasing its zoom level to the highest displayable level. To accommodate the larger space requirement, the zoom levels of its neigbouring regions will be decreased. This style of zooming is similar to those found in TableLens [16] and DateLens [2], where the new space requirements caused by zooming into a cell in a table is satisfied by reducing the sizes of its neigbouring rows and columns. As before, priority to space is given to the more highly intersected regions. In order to convey the transition and to link the two views, changes in zoom levels are animated [19].

3.5 New Queries

MusicLand supports refinement of query with two features: showing a list of selectable attributes for the selected music record, and linking the new queries to their originating queries.

Selectable attributes for the selected music record. Once the user selects a music album, MusicLand displays a list of its attributes that is not already used as query terms in the selected query. This list is used to guide the user in refining their query (i.e., to answer the question, "where *can* I go next?"). Selection of these attributes will initiate another query. For example in Figure 4, once the user has selected the album "Various Artists/Artificial Intelligence", a list of attributes for that album is displayed.

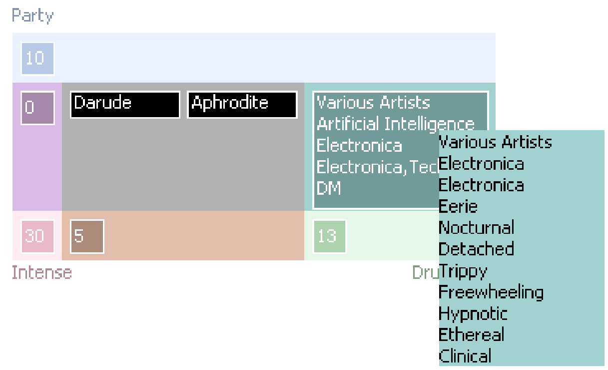

Figure 4. Attribute list for the selected music record.

Linking of new to old 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. For example in Figure 5, selecting "Eerie" from the selectable attribute list (Figure 4) results in a new query, which is linked to the original query by a turquoise coloured line, the same colour as the background of the music record box. The queries therefore constitute a linked list that provides a history of query evolution. This trail may help to give the user a sense of "where have I been" and help him refine his target (i.e., to answer the question "where should I look for next?").

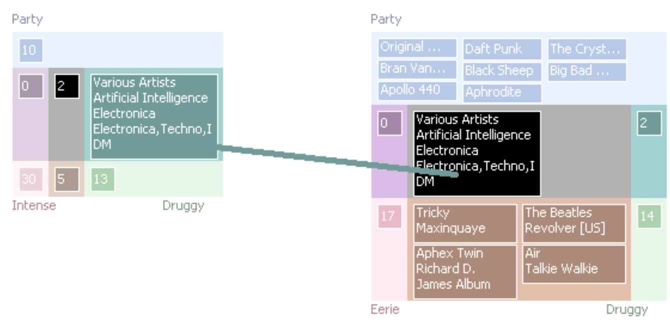

Figure 5. New query.

Basically, queries are displayed in MusicLand in rows and columns to produce a visually clean and orderly display. Another design choice in MusicLand is to allow vertical scrolling only, since horizontal scrolling is believed to be more disturbing than vertical scrolling [6], and 2D scrolling can be very disorienting. As a result of this decision, only rows but not columns are added to the display to make space for new queries. If the originating query is surrounded by existing queries, a new row is inserted between existing query rows to avoid long links. To avoid excessive link crossings, new queries are added either immediately below, or to the sides of the original query (i.e., not added to the diagonal).

For the colour coding of query regions, it is obvious that the display will not be able to encode each query term with a distinctive colour as the user initiates more and more queries. In MusicLand, the three primary colours are reused in new queries. While such an encoding scheme is appropriate for shared search criteria among the queries, it is disturbingly misleading for changed query terms since we perceptually group objects by their colour.

Changes to the original display scheme are therefore needed to address issues in displaying multiple queries: (1) to resolve the unwanted colour grouping, and (2) to avoid large amount of scrolling. MusicLand uses a “time” metaphor to address both issues, where old queries fade in colour and decrease in size with time, similar to physical objects in the world and in our memories.

Colour encoding with new queries: To minimize the unwanted colour grouping effect, regions corresponding to old query terms that are not reused will fade with each new additional query. Eventually, colour of these regions will completely fade away, leaving a more colour saturated trail of reused query regions (e.g., the turquoise trail shown in the left-side of Figure 1). The process is reversible. When the user revisits the query display, and initiates a query using faded terms, the corresponding regions will incrementally regain their original colour.

Space allocation to new queries: In addition to vertical scrolling, MusicLand resizes old queries to maximize information content within the viewable area. Except for the reused query terms, old queries should be less important to the task at hand than new queries. Thus reducing the size of these old query regions has the benefit of freeing up space, *and* focusing the user's attention to more relevant areas of the display. Reducing the size of old query regions thus works in harmony with the "fading colour" idea mentioned above, and these two operations together constitute the “time” metaphor. An example showing this effect is shown in Figure 1, where the “Intense” region in the first query (the bottom left region of the far top-left query display) has faded and reduced to its minimal size. Similar to the colour fading, this process is also reversible. Reselecting (by reclicking) a minimized region will restore it to its maximal display size.

Animation is used to convey the idea that new queries are extensions of their originating queries. More specifically, animation shows,

(1) "Extension" of the originally selected music record into to the new query focal region;

(2) "Growing" of the new query from this focal region outward to the rest of the query display.

3.6 Dataset and Implementation

This section explains the music album dataset used in the current prototype, and highlights the major mechanisms behind MusicLand.

3.6.1 Dataset

MusicLand uses a mp3 database created by Eric Brochu from the University of British Columbia, which was originally collected from the All Music web site (www.allmusic.com). It consists of 8556 mp3 files extracted from 714 albums by 315 artists. The main genres represented in the database are rock/pop and electronica. Each album is labeled with English terms that describe the tone of the album (e.g., “intense”, "aggressive", or "bittersweet") by a software system called MILQ (Music Interpreted as Lexical Qualifier [4]). Information about the music is stored in ASCII files. Here is one sample:

ALB	Fever to Tell
ART	Yeah Yeah Yeahs
REL	Apr 29, 2003
GEN	Rock
STY	Indie Rock, Garage Punk
TON	Cathartic, Exuberant, Boisterous, Passionate, Brittle
PAT	/cs/beta/SCRATCH/music/mp3library/Yeah Yeah Yeahs/Fever to Tell

where,

ALB is the album title,
ART is the artists,
REL is the release date,
GEN is the genres,
STY is the styles,
TON is the tone labeled by MILQ, and
PAT is the file path of the mp3 file on the server computer.

MusicLand uses ART, GEN, STY and TON as potential query terms. Further clustering of retrieved results is based on ART, GEN or STY. MusicLand uses these terms based on user-defined options.

3.6.2 Implementation

The current MusicLand prototype was implemented in Java with the Eclipse IDE using only the swt library. Since most of the data processing is connected to the visualization, a flat architecture instead of a client-server architecture was adopted.

Before queries can be initiated, the music data (as ASCII files) have to be read into memory and analyzed. Once the user enters the query terms, the system searches amongst the loaded data, and slots the matched music records into the different regions of the Venn diagram. The matching can be general (i.e., match any of the attributes in the music data), or selective (match only certain attributes, e.g., ART, GEN, STY, or TON) as defined by the user. This option thus serves as a coarse filter to speed up the retrieval process, especially for large databases. If the user has chosen to further cluster the data into music piles in the display, the system will cluster the retrieved results according to the user defined clustering attribute (e.g., by ART, GEN, or STY).

Now the system is ready to display the retrieved results. It first calculates the amount of space required by each of the region if it were to be display in full details. If the available space cannot accommodate full details for all regions, the system then decreases the zoom levels of the regions until the whole query can fit into the allotted space. The reduction process is done in a sequence based on the relevance of the regions, i.e., highly intersected regions will be given higher space priorities.

The system then looks for the next available slot on the screen to put the new query display. To place the new query as close to the originating one as possible, the system looks at the available empty spaces and select one that is an immediate neighbour of the originating query display, i.e., one that is to the immediate left, right, or below the originating query display. If a suitable spot is not found, a new row is inserted below the originating query since only vertical scrolling is allowed in MusicLand.

Animation is added to preserve object constancy when the user zooms into a region that is previously out of focus, and to link the new query to its originating query visually. In both cases, the final positions of the regions are calculated first. The system renders the displays in increments based on the number of steps in the animation. To ensure the eventual display reaches the calculated positions (and avoid round off errors, for example), the display assumes the calculated positions in the last animation step.

4 Scenario of Use

In the introductory section, I have described a scenario of e-commerce as one of the intended scenario of use for MusicLand. In this session, I will describe a different scenario of use and string the visualizations introduced in Section 3 in a simple walk-through.

Imagine at the end of a very long day, a user wishes to listen to a piece of music that can help him relax and ease into the evening. He thus inputs three query terms into the system: "Soothing", "Peaceful" and "Carefree". The system displays the query results as in Figure 6a. Noticing that none of his record albums satisfy all the query terms, the user explores the search results further by zooming into different regions of the display (Figure 6b). Seeing his favourite band (Thievery Corporation) in the “Peaceful” region, he plays the record by double clicking the record (not implemented). After listening to the record, he wishes to continue his music experience in a similar mood. To explore his options, he selects the “Thievery Corporation” album to which he has just finished listening in order to reveal the attributes of the album (Figure 6c). He then selects “Melancholy” and “Reflective” to initialize another query (Figure 6d). He notices the new query is visually linked to the old by the album he originally selected (i.e., the “Thievery Corporation” album) with a coloured link. 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 (not implemented).

5 Evaluation

A formative evaluation was conducted to assess the strengths and weaknesses of the MusicLand prototype. Two participants, one male and one female, both experienced computer users, were recruited for the study. The study sessions took the form of interviews, and lasted about 30 minutes per session. The session began with a brief explanation of the intended scenario of use for MusicLand, and followed by a demonstration of the prototype.

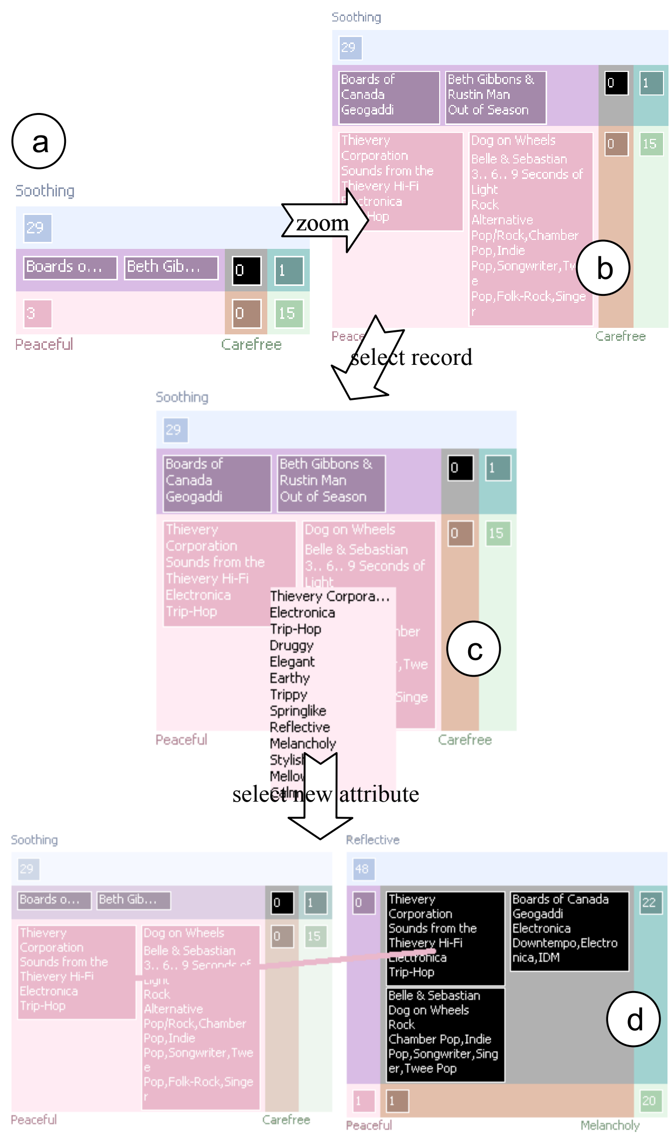

Figure 6. Scenario of use. (a) initial query; (b) zoomed in view of query; (c) display of selectable attributes of the selected music record; (d) display of new query

Participants were then left to explore the interface, and their impressions of MusicLand were solicited. More specifically, the study was conducted to answer the following questions:

(1) Based on the scenario of use, do you feel that MusicLand supports (impedes) exploratory browsing?

(2) What are the main strengths and weaknesses of MusicLand?

5.1 MusicLand Strengths

Both participants reviewed MusicLand favourably, and indicated that the tool would be an improvement over conventional search engine displays for the task of exploratory browsing based on the limited experience they had with it. Neither of them believed MusicLand would impede exploratory browsing.

As to why and how MusicLand may be helpful for exploratory browsing, my study participants commented as follows,

(1) *Supports interpretation of query results in context*. One participant liked the arrangement of retrieved results in context of the query terms. He believed the visual spatial arrangement helped him select the music album of interest when he navigated from region to region within the query display. Another participant liked the overall layout of the query display, since the highly intersected areas are visually salient, thus focusing her attention to the relevant areas of the display.

(2) *Provision of possible query term choices*. The participant said the list of selectable query terms was helpful for her to find music albums that were similar to the ones she liked in the current query.

(3) *Guiding towards future queries*. One participant said the old queries provided a trail that helped him remember his previous queries, and this memory was helpful to define his future queries.

(4) *Fading and desizing of query regions*. One participant commented favourably on the colour fading and desizing of unused query regions. He believed such techniques visually removed unimportant queries without physically removing them, since he may revisit those queries in the future.

Overall, the participants' comments were encouraging. This indicates that MusicLand has largely achieved its design goals to support exploratory browsing by providing context for interpretation, and by guiding future queries with possible query terms and query history.

5.2 MusicLand Weaknesses

The participants also indicated areas of improvements for MusicLand:

(1) *Asymmetry of query term display*. One participant was initially confused about MusicLand's transformed Venn diagram, and commented that such a representation would need to be learned. He especially disliked the asymmetry of the display, where the area and location of the top query term are different from those of the bottom ones, even thought those terms should be conceptually identical.

(2) *Lack of indication of new query terms in new query*. A participant wanted to see the new query terms (the ones that are added to the originating query) displayed explicitly, as an annotation for the query history. Even though it is possible to deduce the new query terms by comparing the terms between the old and new queries, he believed the process would require too much work.

(3) *Lack of Previews/Closing of query displays*. One participant would like to see a preview of the query before committing to it. In a similar line of thought, the other participant wanted to close existing queries. Even though she believed the fading and desizing of query display regions were helpful in visually "removing" unimportant queries, it would be useful to be able to physically remove unwanted queries altogether.

(4) Lack of ability to store favourite albums. One participant wanted to be able to store her favourite albums away from the main display so as to be able to revisit them later in the query process.

(5) Lack of ability to spatially rearrange query displays. A participant wanted to rearrange the query displays to better reflect her mental model of the query results.

(6) Lack of album details. Both participants would like to see the full details of the album when selected. The information is usually published at the back of albums, and is different from the selectable attribute list provided by the music database.

In addition to my participants' comments, I believe MusicLand also has the following weaknesses:

(1) Scalability. Despite my attempts to display a large amount of music record data by semantic zooming and piling of music, MusicLand is still very limited in terms of display capacity. In my testing with 1000 albums, there were times where a single region contained over 50 albums, and could not be displayed even at the "Minimal" zoom level. One possible solution to this problem is to apply a filter to the music database by genre, since a lot of users may prefer only a few types of music. Another approach is to display the result as a list outside of the main display.

(2) Colour encoding. I feel that the perceptual layering in MusicLand is not as effective as it should be. The three main layers are not immediately distinct. The problem may be due to the addition of the music record boxes on top of the regions. Since the music record boxes are of the same colours as the regions, and since it is easier to create a layering effect using the same hue than with different hues, these boxes may introduce a more salient layering effect within the regions than between the regions. This effect can be seen by removing the boxes from the query display (e.g., by comparing Figure 2a or b, to c).

(3) Colour fading. While in theory, colour fading may be a good idea to remove attention from unimportant query regions, it also renders the text illegible, since faded colours provide a smaller contrast to the white text (e.g., Figure 1, bottom-left query region "Intense" of the top-right query display). One possible solution is to dynamically adjust the text of the display based on its background colour.

5.3 Lessons Learned

This project has been a valuable learning experience, where I learned new lessons, and confirmed some of my old thoughts.

(1) *The importance of prototyping*. At the beginning of the project, I spent a lot of time in designing the visualization and interaction of MusicLand without any implementation consideration. That decision was conscious and was based on my prior experiences in interface developments and my

collaboration with an experienced human-computer interface researcher. In this case, I went through four prototypes at the project proposal stage, where many of these prototypes were fairly developed. My last prototype contained enough details that very few design decisions had to be made during the implementation process. This decision turned out to be correct, as the amount of time required to implement MusicLand was much less than I originally anticipated.

(2) *Colour is difficult*. My knowledge in colour perception turns out to be inadequate to realize my vision for MusicLand. Despite having some basic knowledge and harvesting the wisdom of Van Laar (who is an established human factor researcher specializing in using colours in information display), MusicLand still requires work on the use of colour. Also, the current MusicLand does not accommodate users with colour perceptual difficulties.

(3) *Provision of context and details has a price in scalability.* This Focus+Context problem is a very difficult one. While MusicLand may be effective in handling up to 30 unclustered music data records per region, and about 20 queries on the display, it is easy to imagine how awkward it will be to display larger amount of data and queries. It seems to me that there are two distinct approaches in displaying information: overall/thematic (e.g., ThemeScape and ThemeRiver), where thousands of data records can be accommodated, but individual record details cannot be easily shown, or filtered/individualistic (e.g., MusicLand), where only a small subset of the data can be comfortably displayed with some degree of details at the individual record level.

6 FUTURE WORK

Based on the participants' comments, a number of improvements can be made to the current MusicLand prototype:

(1) Add additional query terms to the links to annotate the query trail;

(2) Display full record details apart from the main display;

(3) Allow saving of favourite albums;

(4) Consult a colour specialist to improve on the perceptual layering and colour encoding.

Further development of MusicLand involves launching and playing the music albums. The reason why this was not included in the current project is because of the care and consideration such an operation requires. For example, if the user merely wishes to sample the album, then only the tracks that are most representative of the corresponding query terms should be played. For those tracks, perhaps only the most "interesting" part should be played (e.g., by extracting the chorus sections of each song as described in [9]). However, if the user wishes to listen to the album, then the whole album should be played. Since these considerations are beyond the scope of information visualization, I have decided to exclude it from the current project.

7 CONCLUSION

This paper introduced MusicLand, a tool that is designed to explicitly support exploratory browsing with music album data. In contrast to direct search, the user in exploratory browsing may not have a clear and precise idea about his target, which may be refined or completely changed over the course of his quest with multiple queries. MusicLand explicitly supports such activities by visually clustering retrieved music album data using a rectangularized Venn diagram to allow interpretation of the retrieved results based on the user's query terms. This spatial presentation is reinforced by perceptual layering and colour encoding, where the most intersected and relevant areas are placed on higher visual layers, and intersected areas are encoded with a mixture of the corresponding primary query region colours.

By displaying a list of possible query terms based on the user's current album selection, MusicLand guides the user in query formulation and target refinements. Query history is preserved in MusicLand as a visual reminder of query evolution. Using the "time" metaphor, old and unused query regions fade and desize to avoid attracting users' attention. The overall effect is to a visually salient trail of query linked by their common query terms.

Semantic zooming with four zoom levels (Full, Partial, Minimal, and Nil) and piling of music album data are used to maximize the display capacity of MusicLand. Users can select any region within the display to increase its zoom level, and the transition is animated to link the two zoom views.

Participant opinions and impressions solicited with a formative evaluation showed encouraging responses, perhaps indicating the viability of MusicLand's design ideas. It would be interesting to continue developing MusicLand in the music domain, and/or apply its design principles to other domains (e.g., digital library and e-commerce) to further study the exploratory browsing behaviour, and to investigate potential roles computer tools can play in such activities.

REFERENCES

[1] C. Ahlberg, and B. Shneiderman. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. *Proceedings of SIGCHI '94*, 313-317, 1994.

[2] B, B. Bederson, A. D. Clamage, M. P. Czerwinski, G. R. Robertson. DateLens: A Fisheye Calendar Interface for PDAs. *ACM Transactions on Computer-Human Interaction*, 11 (1):90-119, 2004.

[3] E, Berenci, C. Carpineto, V. Giannini, and S. Mizzaro. Effectiveness of keyword-based display and selection of retrieval results for interactive searches. Lecture Notes In Computer Science. *Proceedings of the Third European Conference on Research and Advanced Technology for Digital Libraries*, 106 – 125, 1999.

[4] E. Brochu. MILQ. *Master of Science Thesis at the University of British Columbia*, 2004.

[5] M. Chalmers. Using a landscape metaphor to represent a corpus of documents. In A. U. Frank and I. Campari, editors, Proceedings of the European Conference on Spatial Information Theory (COSIT '93), volume 716 of Lecture Notes in Computer Science, pp. 377-390. Springer, 1993.

[6] R. L. Duchinichy, and J. Kwahk, J. Readability of text scrolled on visual display terminals as a function of window size. *Human Factors*, 25:683-92, 1983.

[7] S. Dumais, E. Cutrell, J. J. Cadiz, G. Jancke , R. Sarin , D. C. Robbins. Stuff I've seen: a system for personal information retrieval and re-use. *In Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, 72-79, 1993.

[8] C. Fluit, M. Sabou, and F. van Harmelen. Supporting User Tasks through Visualisation - Of Light-Weight Ontologies. S. Staab and R. Studer ed. *Handbook on Ontologies in Information System,* Springer-Verlag, 1999.

[9] M. Goto. SmartMusicKIOSK: music listening station with chorus-search function. *Proceedings of the 16th annual ACM symposium on User interface software and technology*, 31-40, 2003.

[10] S. Havre, E. Hetzler, K. Perrine, E. Jurrus, N. Miller. Interactive visualization of multiple query results. *Proceedings of IEEE Information Visualization Symposium 2001*, 105-112, 2001.

[11] S. Havre, E. Hetzler, P. Whitney, and L. Nowell. ThemeRiver: Visualizing Thematic Changes in Large Document Collections. *IEEE Transactions on Visualization and Computer Graphics* 8(1), 9-20, 2002.

[12] M. A. Hearst. Using categories to provide context for full-text retrieval results. In *Proceedings of the RIAO '94, Intelligent Multimedia Information Retrieval Systems and Management*, 115-130, 1994.

[13] 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)L:131-161, 1996.

[14] A. Leuski, and J. Allan. Lighthouse: Showing the way to relevant information. *Proceedings of IEEE Information Visualization Symposium 2000*, 125-129, 2000.

[15] X. Lin. Visualizing for the Document Space. *Proceedings of IEEE Visualization '92*, 274-281, 1992.

[16] 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.

[17] P. Ravasio, L. Vukelja, G. Rivera, and M. C. Norrie. Project infospace: From information managing to information representation. *Proceedings of Interact 2003---Ninth IFIP TC13 International Conference on Human-Computer Interaction*, 2003.

[18] M. Ringel, E. Cutrell, S. Dumais, and E. Horvitz, E. Milestones in time: the value of landmarks in retrieving information from personal stores. In *Proceedings of Interact 2003*, 184-191, 2003.

[19] G. G. Robertson, S. K. Card, S.K., and J. D. Mackinlay. The Cognitive Coprocessor Architecture for Interactive User Interfaces. *In Proceedings of Symposum of User Interface Software and Technology*, 10-18, 1989.

[20] A. Spoerri. InfoCrystal: A Visual Tool for Information Retrieval and Management. *Proceeding VIS 93*, 11-20, 1993.

[21] J. A. Wise, J. J. Thomas, K. Pennock, D. Lantrip, M. Pottier, A. Schur, and V. Crow. Visualizing the Non-Visual: Spatial analysis and Interaction with Information from Text Documents. *Proceedings of IEEE Information Visualization Symposium '95*, 51-58, 1995.

[22] A. Woodruff, R. Rosenholtz, J. Morrison, A. Faulring, and P. Pirolli. A comparison on the use of text summaries, plain thumbnails, and enhanced thumbnails for web search tasks. *JASIST* 53(2):172-185, 2002

[23] T, Uphill. Consistency, clarity and control: development of a new approach to WWW image retrieval. *Bachelor of Information Technology Thesis at Australian National University,* 2000.

[24] D. L. Van Laar. Psychological and cartographic principles for the production of visual layering effects in computer displays. *Displays*; 22:125-135, 2001.