Visual Navigation within Conceptual Spaces

Hans Jörg Müller*
Universität Münster
joerg.mueller@uni-muenster.de

Antonio Krüger*
Universität Münster
antonio.krueger@uni-muenster.de

Abstract

Finding objects like music albums or wines in large databases is an inherently difficult problem, especially if users do not know exactly what they are looking for. We propose the use of Conceptual Spaces to lay out items spatially and provide a set of techniques for Zoomable User Interfaces to enable navigation within these spaces. Conceptual Spaces place items along meaningful dimensions, so users can find the items they are looking for, and similar items are placed near to each other. We implemented a prototype to enable navigation within a collection of music albums, and believe that the concept has the potential to generalize to all kind of real world items like pictures, food, wines or books that can be described by perceptual qualities as time, color, emotion or taste.

1 Introduction

Imagine a user entering her favorite music store. She does not know exactly what she is looking for, she already knows some of the albums available, and the albums are laid out according to some schema. If the store is good, similar albums are placed near to each other, and if she understands how the albums are ordered, she can find the albums she is looking for. But there are many dimensions the albums could be ordered by: genre, artist, release date, region, etc. Albums could also be similar according to these different dimensions, they could be from similar genre, similar time, or similar region. When the collection is digitally available, e.g. as her personal mp3 collection or at an online music store, the spatial arrangement of the items could be adjusted according to her current search criteria. Both in the real world as in the digital case, at any point in time there should be some way for her to know where to find each album. We propose to layout the albums according to Conceptual Spaces and suggest a navigation method within this representation with a Zoomable User Interface. This approach scales both to different display sizes, so the same interface can be used on a mobile mp3 player, a desktop PC or a wall-size display, as to different domains, like picture galleries, restaurants, wine stores, book stores etc. In the following we will explain our ideas by first reviewing relevant related work in section 2, where similar approaches and the basic foundations of Conceptual Spaces are discussed. In section 3 we will provide a description of how multiple views on data sets can be organized by Conceptual Spaces. We will also introduce visual techniques in the context of Zoomable User Interfaces that enable users to visually navigate through Conceptual Spaces and thus

*Institut für Geoinformatik, D-48149 Münster, Germany

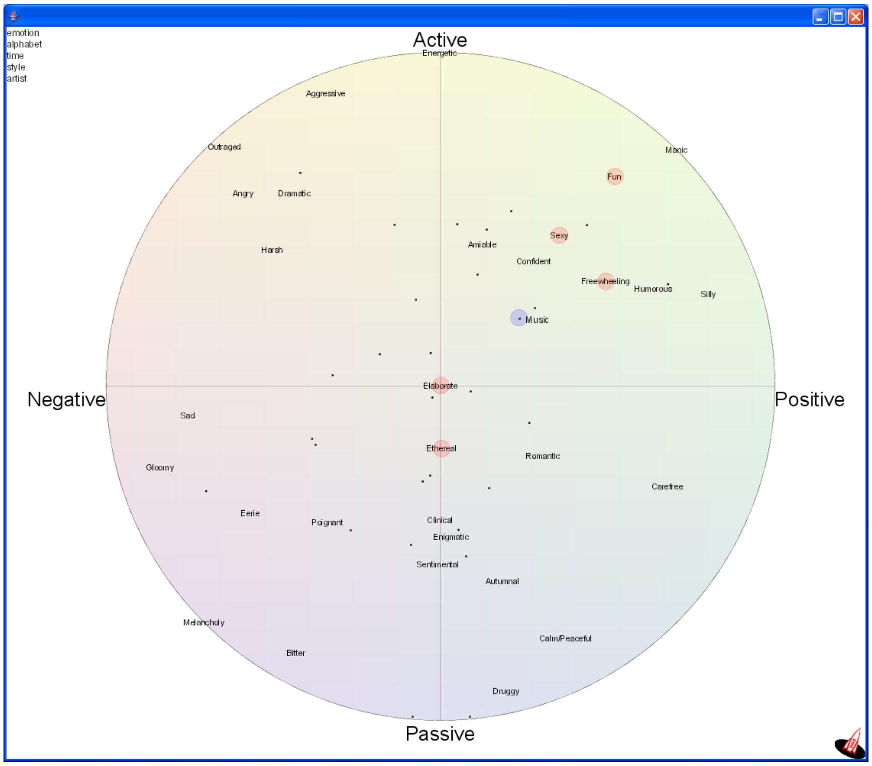

Figure 1: Searching for happy music albums in the emotion perspective. The big labels show landmarks, the small dots are individual albums. For the selected album 'Music' from Madonna, corresponding landmarks like 'Fun' are highlighted.

helping them to discover similar objects in a given data set. We will present a prototypical implementation of our concepts in the domain of music albums that can be visually explored along the dimensions of time, geographic space, and emotion.

2 Related Work

Shneiderman [Shn96] introduced the "Visual Information Seeking Mantra" of "Overview first, zoom and filter, then details-on-demand". But there is still the question of how items should be ordered spatially. To better understand the benefits of our idea to use Conceptual Spaces for this layout task, we will first review related approaches and systems, and then provide a brief introduction into the basics of Conceptual Spaces.

2.1 Similar Approaches

Several approaches have been suggested for the layout of items along semantically meaningful dimensions. Mackinlay [Mac86] showed how graphical presentations of relational information could be generated automatically and how items could automatically be laid out along meaningful dimensions in scatter plots. His focus was not to help the user finding specific items, but to compare items according to relational properties, so there was no work on navigation and he used static representations. Film finder [JS94] provides navigation within a movie database which is laid out in a starfield display with year of release on one axis and a popularity score on the other axis. Categories are color coded, and it is possible to filter the items by title, director, length etc. using slider bars. Zooming is possible using the newly introduced zoom bar. In contrast to these approaches, we propose not to show independent domains like price and time at the same time, but to provide different perspectives between which the user can switch. So, the structure of each domain can be maintained. The Dynamic Home Finder [AS94] enables visual information seeking within a database of houses that are for sale. The houses are laid out on a map, and constraints on further dimensions like price can be set with sliders. Life streams [FFG96] provide a personal diary for all digital items like documents, emails, and pictures. Because navigation is only possible along the time dimension, file names and hierarchical folders are rendered obsolete. We generalize this approach to include other dimensions than time and propose Zoomable User Interfaces for navigation. Semantic regions [Kan03] are rectangular regions on a plane with specific semantics. Items are displayed in their corresponding regions and can be displayed in multiple regions at one time. A map of the US could be divided into multiple regions, for example. Semantic Regions are meant to represent the mental models of the user. By generalizing Semantic Regions to Conceptual Spaces, we provide not only a finite number of discrete regions, but continuous dimensions to order the items by. Visual information seeking tools have to provide both focus and context information, that is provide detailed information for some item while helping the user to maintain a sense for its context. This can be solved either by providing separate detail and overview windows, by distorting the space with a fisheye or hyperbolic [LRP95] lens or by zooming. Because it is so important that the user understands the structure of the space, focus+context techniques that distort the representation of the space on the screen by using fisheye or hyperbolic views are less suited for Conceptual Spaces. Zoomable User Interfaces (ZUI) [PF93] enable smooth zooming in 2D and can provide focus information when zoomed in and context when zoomed out without distorting the space. Dynamic Query Filters [Shn94] allow the user to interactively refine a query to a database and view the results immediately, either by adjusting sliders or by drawing selection boxes around the items. Halos [BR03] are circles drawn around off-screen items that are just large enough to reach into the border region of the display, so the user can infer the off-screen location of the item from the curvature of the fraction of the circle that is visible on the display. In section 3 we will show how we make use of some of these techniques for the purpose of visual navigation in Conceptual Spaces. Beforehand we will discuss some of the basics of Conceptual Spaces.

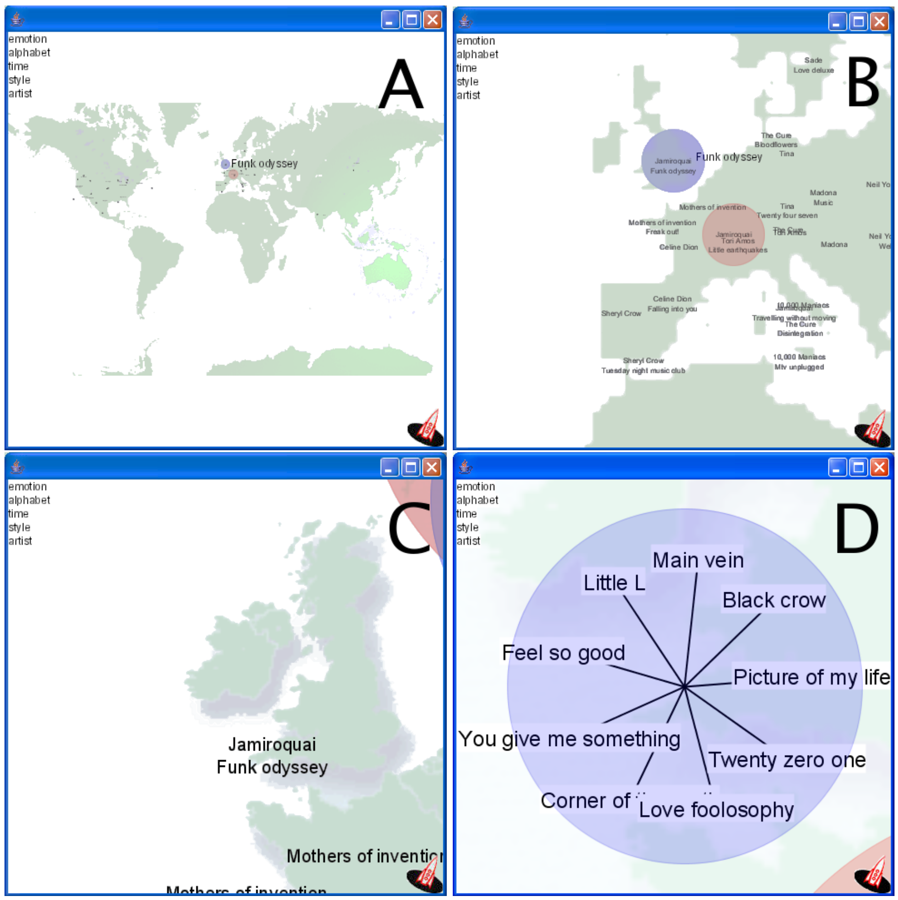

Figure 2: Zooming to an album of Jamiroquai in the map perspective (A, B, C), and playing songs from that album (D).

2.2 Conceptual Spaces

A unifying model of how to describe properties of items spatially and thus how to lay them out so that users can find the items they are looking for according to different properties is described by Conceptual Spaces [Gae00]. Conceptual Spaces are a framework in the family of geometric models to describe conceptual knowledge using vector spaces. A Conceptual Space could for example consist of the Domains time and geographic space, and a number of appointments could be points within this space. Objects are represented as points in the space and Properties as regions in the Domains, like "today" or "at home". This model emphasizes similarity as a function of distance in the Domains, so a house would be geographically similar to the neighbors house because in geographic space they are near to each other. In Conceptual Spaces, a Quality Dimension is a dimension with ordered values, like time. A set of Quality Dimensions that belong together, like longitude and latitude, are called integral, and independent Quality Dimensions, like latitude and time, are called separable. A Domain is defined as a set of integral dimensions that are separable from all other dimensions, like longitude and latitude for geographic space. A Conceptual Space is then a collection of Domains, like space-time. Domains can describe all kinds of perceptual qualities, like color, taste, geographic space, time, or emotion.

3 Navigation within Conceptual Spaces

If we want to provide navigation within a collection of items, each item can be represented as a point at a specific location in an adequate Conceptual Space. If items are laid out according to the Conceptual Spaces of the user, the interface to navigate within these spaces has to be designed to support the structure of Conceptual Spaces. Because the location of each item describes its properties, the user should always be aware of the relative position of all relevant items in the space.

3.1 The ConceptSpace Application

As an application domain, navigation within a collection of music albums was chosen. For this purpose, the albums are laid out spatially in Conceptual Space Domains. In the time Domain, for example, they are ordered by their release date. Thus, the user can browse the space to find music albums that fit her search criteria, for example "Music from the 80's" (Figure 4). We started considering 1D and 2D Domains only, leaving the possibility to extend the work to 3D later. We implemented three interval scaled spaces, 'emotion', 'time', and 'artist/map', and two nominal scaled spaces, 'alphabet' and 'style'. Because Conceptual Spaces are best suited for interval and ratio data, we present the corresponding perspectives in detail. To deal with the possibly infinite size and resolution of Conceptual Space Domains and to provide both focus and context information, we used a Zoomable User Interface. ZUIs allow for both infinite screen size and infinite resolution by providing continuous zooming in 2D. In the Domain of geographic space (Figure 2), for example, an overview of the whole world can be presented at a low zoom factor (A), and the user can continuously zoom into a certain point to get more details (B, C) and see the individual songs in the album (D).

3.2 Information Scent

If all the information items are laid out in a 2D space, the user has to know where to look to find the item she is looking for. This can be compared to finding a trace that leads to the prey while the user is hunting. In principle, there is a trace available that leads to the information item the user is hunting, she only has to find it. In the terminology of Information Foraging [PC99], there has to be some information scent for each item that tells the user where to look. This is exactly what Conceptual Spaces do: If the user understands the structure of the space, like Ekman's emotion circle [Ekm57], that categorizes emotions along an active-passive and a positive-negative scale (Figure 1), she knows where to look to find music

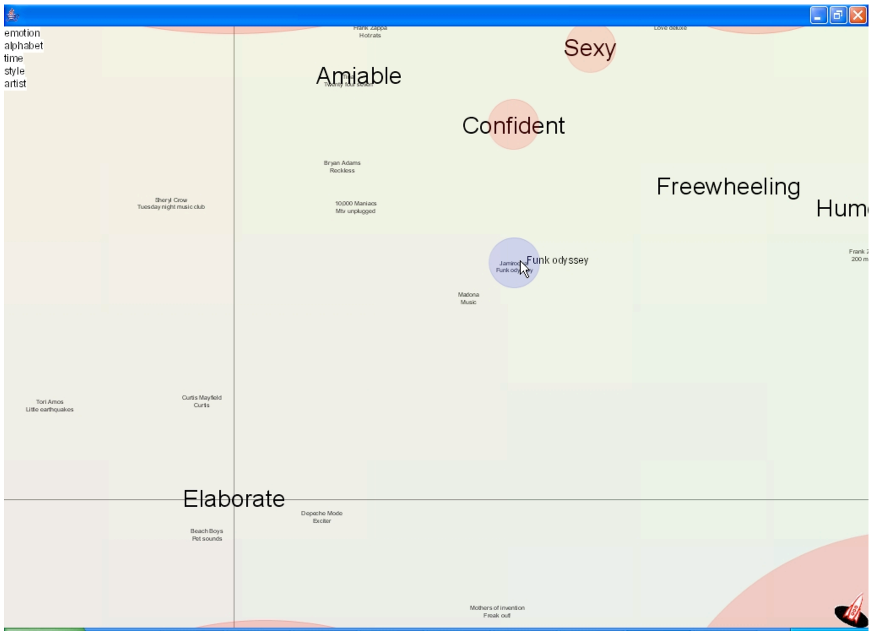

Figure 3: Searching for albums similar to “Funk odyssey” in the emotion perspective.

albums that contain the emotions “fun” or “fear”. If she does not understand the structure of the space, however, she does not know where to look at all. So it is very important to provide representations for each domain that are understood by the user. This can be achieved by providing multiple representations for each domain, like multiple models for emotions, matching the individual representations of the user, by trying to make the representations self-explaining using appropriate landmarks, or by explaining the representations to the user. In addition to the structure of the space a number of additional techniques can be used to provide further information scent. For example, the representation of the item itself on the screen can provide further information about the item, either by displaying text or icons. The technique of Semantic Zooming [PF93] changes this representation depending on the zoom factor (Figure 2), and represents a music album as a dot at a low zoom factor (A), as a string containing the title and artist at an intermediate zoom factor (B and C), and as a list of songs at a high zoom factor (D). Tooltips provide additional “scent on demand” when the pointer hovers over an item (Figure 1). For selected items that are currently not on the screen, Halos are used to highlight them (Figure 3).

3.3 Using Multiple Perspectives

For each kind of items, a variety of Domains can be used. For our example of music albums, the release date (time), the city of recording (geographical space) or the emotion that is associated with the album are important. For wines, the year of production (time), the region

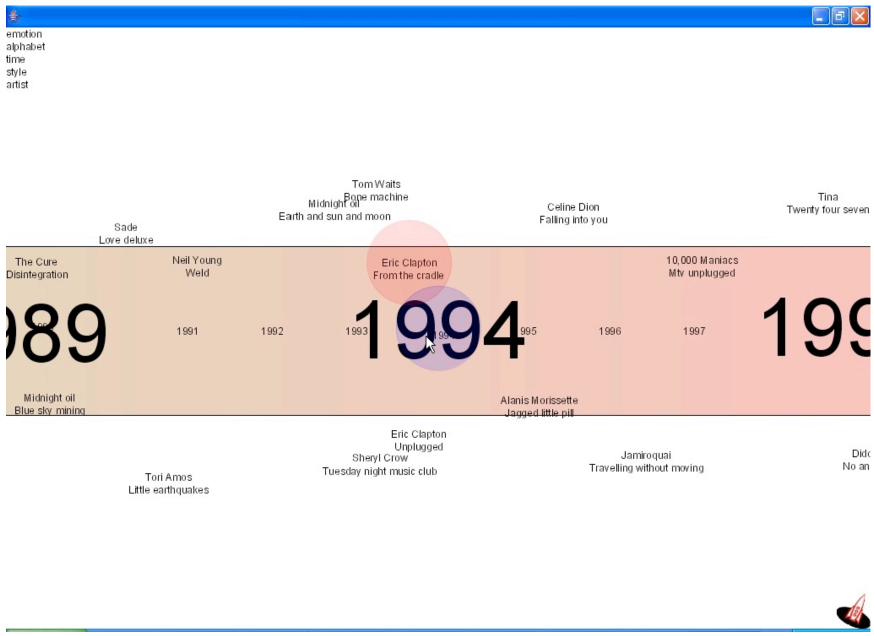

Figure 4: Browsing for albums from the 90's in the time perspective.

of origin, the color, tastes and fragrances could be used. The choice of a Domain does not determine the structure of the Domain. Let us look at two examples, time and color. Time, for example, can be visualized as a straight timeline, running from left to right, as a spiral or helix with one round representing one day, one week or one year, as a table with the columns Monday to Sunday, or a number of other representations. The same holds true for color: Possible structures are the color circle, the NCS color spindle etc. Structures for the visualization of various domains have been extensively studied and can be found in [Tuf90], for example. Each representation emphasizes different aspects of the Domain. Because in Conceptual Spaces similarity is a function of distance, each of these representations focuses on a different notion of similarity. In the linear timeline, for example, "today, 8 am" would be nearer (more similar) to "today, 8 pm" than to "tomorrow, 8 am", whereas in a helix with a frequency of one day it could be the other way around. The optimal structure to be used for a Domain depends both on the kind of items that are displayed and on the individual preferences of the user. Each Domain consists of a number of integral Dimensions that are separable from all other Dimensions. There are various possible approaches to deal with the integration of different Domains. One possibility would be to try to integrate several separable Domains, for example, longitude, latitude and time, into a 3D representation. The drawbacks are that this is limited to a total of three dimensions and that the display of separable Domains in a single space is usually not very intuitive. If time and space are integrated in one presentation, Münster today could be as distant and thus as similar to Venice today as Münster today would be to Münster 30 years ago. The meaning of

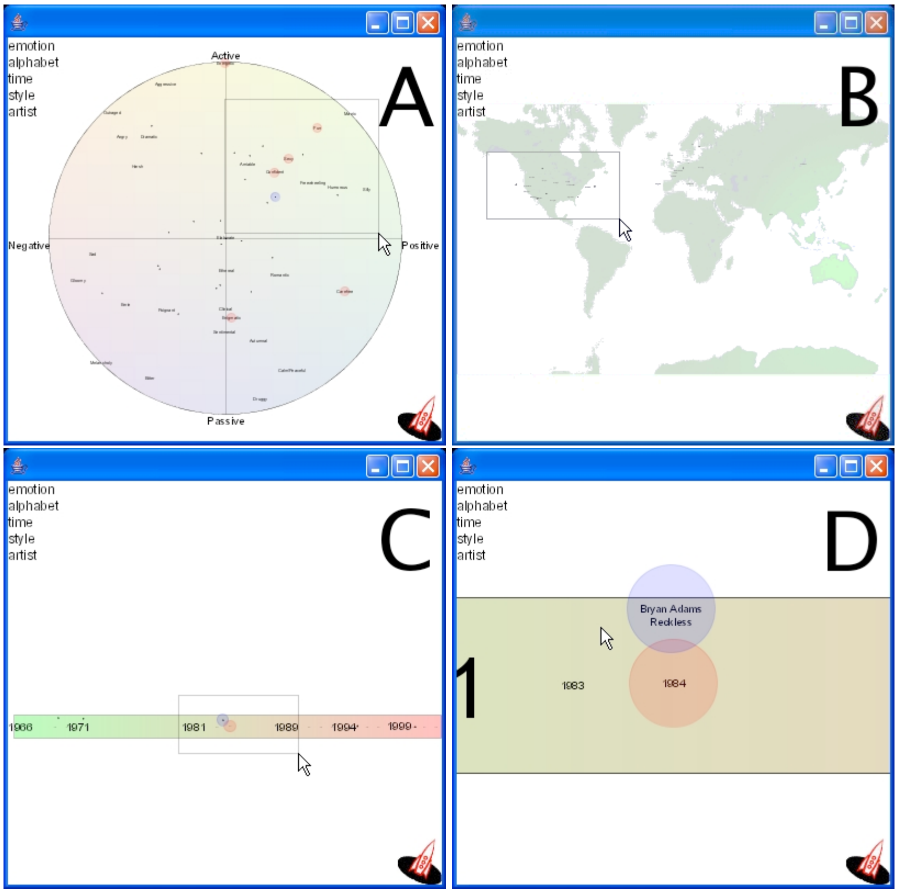

Figure 5: Selecting some happy (A) american (B) music from the 80's (C,D).

this similarity measure is not obvious. When an additional feature, like color or texture is used to present another Domain, there is no intuitive way to navigate along this property. Therefore, we chose to only display sets of integral Quality Dimensions at a time and have a facility to switch between Domains. To maintain the context when switching between different perspectives we implemented an embedded spring layout algorithm to smoothly move the items to their new positions in the next Domain. Thus, both the movement of a few highlighted items and an overview of the general dynamics can be perceived by the user.

3.4 Navigation

ZUIs can provide both panning and zooming for navigation. Because one of our targeted devices has a pen interface and thus provides only a few gestures like click or click-and-drag, we tried to get along providing only zooming and no panning for interaction. Panning could be easily implemented by key presses or combo gestures. Since we obtained satisfactory results without panning, there was no urgent need to introduce panning into the current version of the interface. Usually, without proper navigation tools, users easily get lost in ZUIs by zooming in or out too far. To overcome this problem we provide a semiautomatic navigation mode that pre-selects certain views on the music albums and titles. In any ZUI application, there can be identified some more interesting and some less interesting views. A view that does not contain any items and thus provides no information scent at all is commonly called a desert fog view [JF98]. Desert fog views are especially uninteresting and thus should be avoided. On the other hand, the smallest view that displays all available items, which we call the root view, is of special interest. Furthermore, for each item one view that shows title and artist at a sufficient resolution and one view that shows the songs at a good resolution might also be interesting. We placed a magnetic zone around each of these interesting views, such that the user is automatically zoomed to the interesting views when he enters the magnetic zone. This is similar to the Lodestones and Leylines technique [Jul02] that only allows navigation between interesting views, but is less constraining. For zooming out we used the Leylines technique such that the system always zooms out in the direction of the root view, independent of pointer position. In addition to background elements to visualize the structure of the Domain, like a ruler bar for the time Domain, a number of Landmarks were used to facilitate orientation, like certain time points. We also implemented a 'rocket' button that zooms out to the root view, and back to the previous view on release, to provide a canonical view on the context.

3.5 Filtering

Using zooming alone, the user can only navigate in one perspective at a time. It would be quite difficult for her to find albums that match criteria from multiple perspectives, like "happy European music from the 80's". This combination of perspectives is achieved using Dynamic Query Filters as introduced in [Shn94]. Filters are constructed by drawing selection boxes in different Domains. In contrast to previous approaches, where filters were set using sliders or drop-down list, we can benefit from the properties of Conceptual Spaces. A normal filter would constrain the items displayed by a number of properties they must fulfill. Properties in Conceptual Spaces are defined as regions in Domains, and Natural Properties are even defined as convex regions in Domains. Natural Properties have the characteristic that when two objects A and C have a Property, and an object B is between A and C, and is thus more similar to A than C is to A, and more similar to C than A is to C in respect to the current Domain, then object B also has the Property. Thus, by allowing users to select items from a convex region within a perspective, we allow them to choose a Natural Property from the corresponding Domain. Up to now, only rectangular selection boxes are available in the prototype, but they seem to be sufficient for most cases. If the user draws a selection box in a Domain, all items not contained in the selection are faded, and the user can switch the perspective to refine the selection, thus subsequent filters are always combined by an AND operator. Other Boolean operators like OR and NOT are not supported yet. To remove the created filter, the user has to draw an empty selection box that contains no items, and all items are displayed at their corresponding positions again. For example, to find a music album that has the Properties "happy", "from the United States", and "from the 80's" (Figure 5), the user would change to the emotion perspective, draw a selection box around the "happy" region to select those albums (A), change to the geographic perspective to draw a selection box around the U.S. (B) and then change to the time perspective to draw a selection box around the 80's or simply zoom into this region (C,D). Once she has found the desired album, she would draw an empty selection box to see all albums again.

3.6 Annotating Metadata

To be able to choose the correct place for each item in each Domain, we need a lot of metadata for each item. These can either be generated automatically or annotated manually. The best way would be to generate the metadata automatically when the data item itself is created. For a photo for example, a timestamp and current GPS coordinates and orientation could be annotated. If this is not possible, some metadata could be generated by classifying the data automatically, like the mood for a song or predominant colors for a picture, if such classifiers exist. If the metadata has to be generated manually, it could either be generated by the creator of the data item, like the taste or fragrance of a wine, shared by a community of users over the Internet, as MoodLogic [1] does for music, or annotated individually. Another interesting possibility would be to create online games that create the metadata as a side effect [vA06]. For some cases all users would agree on a categorization, like for the release date of an album, while for the emotions there could be individual differences. In these cases it would be reasonable to make a first guess with a generic method and let the user correct it individually. Conceptual Spaces make it especially easy to annotate metadata by just dragging and dropping the item into the correct region. Furthermore, it does not harm if in some perspectives metadata is only available for some items, because the others can just be placed in some 'no metadata available' region.

4 Implementation

The prototype was implemented in Java using the Piccolo ZUI toolkit [BGM04]. Piccolo provides smooth zooming even with a large number of objects and is simple to extend for advanced ZUI techniques, like semantic zooming, Lodestones and Leylines and Halos. To enable easy applicability to different application areas, like music, pictures, wines etc., the application is designed in a modular way. All functionality that is needed for any application area is implemented in the ConceptSpace framework (Figure 6). The prototype application for music albums, MusicSpace, uses this framework and extends it by functionality that is specific for music albums. The display and layout of the information items are managed by the ZoomGraphLayer. To provide a flexible use of different perspectives within each application, a plug-in structure for perspectives was implemented. Each application can register the perspectives it needs, possibly individually for each users preferences.

[1] http://www.moodlogic.com

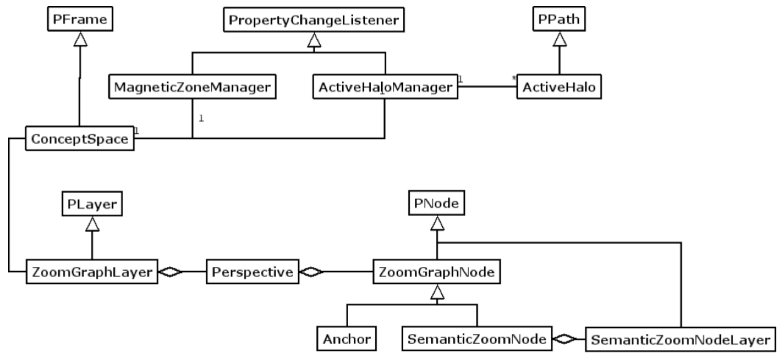

Figure 6: Architecture of the ConceptSpace Prototype.

To provide a smooth transition between different perspectives and to enable the user to compare where highlighted items are in the old and in the new perspective, an adaptation of the embedded spring layout algorithm [Ead84] was used. When the user changes the perspective, invisible "springs" are created between each item and its new position, the "anchor", so that the item is smoothly drawn to its new position within approximately one second. A side effect of this layout algorithm is that overlapping of items is mostly inhibited, but in dense regions items may be not accurately placed at their correct positions at a high zoom factor. Each application can choose how to present its items at different zoom levels. In MusicSpace, for example, a representation as a dot at a low zoom factor, as a text label containing artist and album name at an intermediate zoom factor and as a list of songs at a high zoom factor was chosen (Figure 2). For wines, representation as a dot, a text label with the name and a picture of the bottle could be chosen. For playing the mp3 files when the user clicks on a song, JavaLayer [2] is used. When individual files, like mp3 files or pictures, would be presented in the application, one should consider storing the metadata used by ConceptSpace directly in the files, e.g. in the ID3 tag or EXIM header, but currently we store the metadata in a database.

5 Informal User Trial

We tested the application in an informal user study with 4 members of our lab using a projector and a presentation trackball as interface devices. The first observation was that for all subjects the emotion domain had to be explained but was then used effectively. The first task was to find a happy music album. Starting from the time perspective, subject 1 changed to the emotion perspective, chose one album from the positive active region and zoomed to it (Figure 7 A). The second task was to find an album from Europe from the 80's. The subject changed to the geographic perspective, drew a selection box around Europe to select only albums from that region (B), changed to the time perspective and directly chose

[2] http://javalayer.sourceforge.net

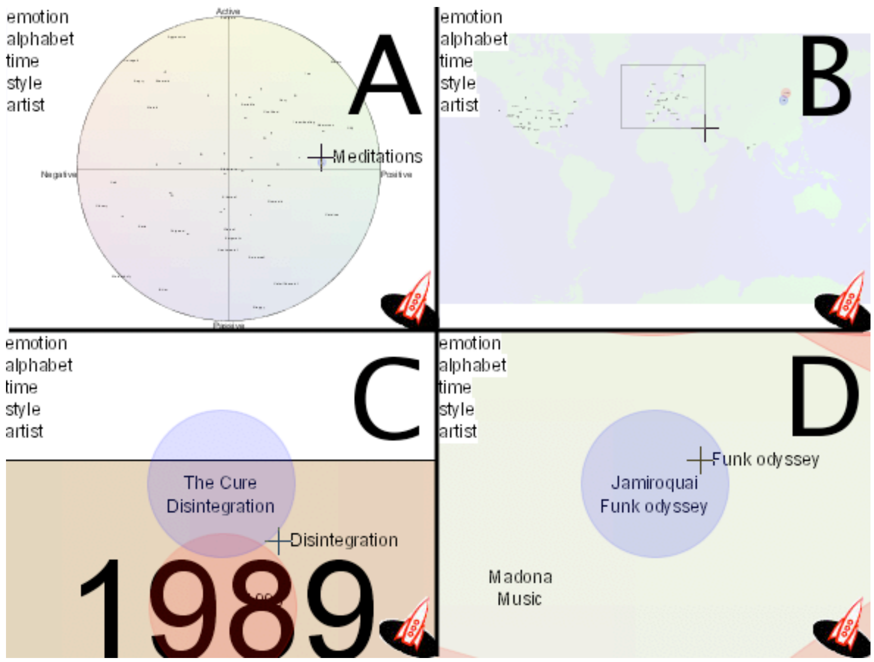

Figure 7: Screenshots from subject 1 during the user trial, while looking for happy music (A), selecting european albums (B) from the 80's (C), and looking for albums similar to "Music" from Madonna (D).

one album from the 80's (C). The third task was less constrained. The subjects were asked to find some album that is similar to the album "Music" from Madonna. All four subjects finally chose to compare the albums in the emotion perspective. The first subject changed to the alphabetical perspective to find the album "Music" and zoomed to it. He filtered by only this one item, changed to the emotion perspective to find out where this item was in this perspective. Having found it, he removed the filter and looked for other albums in the same region. When he found one, he zoomed to it (D).
All subjects could use the interface after a short explanation and considered it useful. The subjects reported they would like to use the interface for a variety of tasks, like music, photos, books and even stamps, and reported that it was fun to use.

6 Conclusion and future work

In this paper, we proposed to layout items according to Conceptual Spaces to provide navigability and suggested a navigation method within this representation with a Zoomable User Interface. We demonstrated a proof-of-concept prototype to show the feasibility of the approach. Since the first user tests described in section 5 are promising, we are are planning a more formal evaluation, i.e. to obtain qualitative measurements of the advantage provided by Conceptual Spaces in visual search tasks. Additionally, to gain more experience with this type of user interface we are investigating other domains than music, such as pictures, wines and books. Considering the fast growing amount of digital media items like photos or songs many people now own and the demand to access these on mobile devices with small screens, we think there is an urgent need to replace hierarchical folders based ordering techniques. Text based search has proven its ability to scale to huge amounts of data, but is only helpful when you already know what you are looking for. We tested the application with about hundred items, but the scalability of our method only depends on how accurately items can be located in the spaces. Emotion, for example, would not scale very well because people would not know where to look exactly. Geographic space, on the other hand, scales very well because items can be located in certain streets or even houses. It would be very interesting to integrate Conceptual Spaces, that work well with ratio scaled data, with techniques like tagging, that work great with nominal scaled data. Finally, we believe that Conceptual Spaces provide a sophisticated conceptual framework for application designers to order items spatially and thus enable visual navigability using multiple perspectives.

Acknowledgements

This work was done at the Tangible User Interface group of Sony Advanced Technology Center, Stuttgart. We would like to thank Martin Emele, Stefan Rapp, Georg Michelitsch, Erica Melis, Anthony Jameson, and the anonymous reviewers for helpful comments and suggestions.

References

[AS94] Christopher Ahlberg and Ben Shneiderman. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Human Factors in Computing Systems. Conference Proceedings CHI'94*, pages 313–317, 1994.

[BGM04] Benjamin B. Bederson, Jesse Grosjean, and Jon Meyer. Toolkit design for interactive structured graphics. *IEEE Trans. Softw. Eng.*, 30(8):535–546, August 2004.

[BR03] P. Baudisch and R. Rosenholtz. Halo: a technique for visualizing off-screen locations. In *Proceedings of CHI 2003*, pages 481–488, Fort Lauderdale, FL, April 2003.

[Ead84] P. A. Eades. A heuristic for graph drawing. In *Congressus Numerantium*, volume 42, pages 149–160, 1984.

[Ekm57] P. Ekman. A methodological discussion of nonverbal behavior. *Journal of Psychology*, 43:141–149, 1957.

[FFG96] S. Fertig, E. Freeman, and D. Gelernter. Lifestreams: An alternative to the desktop metaphor. In *Proceedings of ACM SIGCHI 96 Human Factors in Computing Systems*, 1996.

[Gae00] Peter Gaerdenfors. *Conceptual Spaces: The Geometry of Thought*. The MIT Press, March 2000.

[JF98] Susanne Jul and George W. Furnas. Critical zones in desert fog: Aids to multiscale navigation. In *ACM Symposium on User Interface Software and Technology*, pages 97–106, 1998.

[JS94] Ninad Jog and Ben Shneiderman. Starfield information visualization with interactive smooth zooming. Technical report, College Park MD 20742-3255 USA, 1994.

[Jul02] S. Jul. Predictive targeted movement in electronic spaces. In *CHI 2002 Extended Abstracts*, pages 626–627, 2002.

[Kan03] H. Kang. Personal media exploration with semantic regions. In *Proceedings of CHI 2003*, New York, NY, 2003. ACM Press.

[LRP95] John Lamping, Ramana Rao, and Peter Pirolli. A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proc. ACM Conf. Human Factors in Computing Systems, CHI*, pages 401–408. ACM, 1995.

[Mac86] J. D. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5:110–141, 1986.

[PC99] P. Pirolli and S. Card. Information foraging. *Psychology Review*, 106(4):643–675, 1999.

[PF93] Ken Perlin and David Fox. Pad: an alternative approach to the computer interface. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 57–64, New York, NY, USA, 1993. ACM Press.

[Shn94] Ben Shneiderman. Dynamic queries for visual information seeking. *IEEE Softw.*, 11(6):70–77, 1994.

[Shn96] Ben Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *VL '96: Proceedings of the 1996 IEEE Symposium on Visual Languages*. IEEE Computer Society, 1996.

[Tuf90] Edward R. Tufte. *Envisioning Information*. Graphics Press, May 1990.

[vA06] Luis von Ahn. Games with a purpose. *Computer*, 39(6):92–94, 2006.