



Exploring personal media: A spatial interface supporting user-defined semantic regions

Hyunmo Kang*, Ben Shneiderman

Department of Computer Science, Human–Computer Interaction Laboratory, Institute for Advanced Computer Studies & Institute for Systems Research, University of Maryland at College Park, College Park, MD 20742, USA

Received 4 August 2004; received in revised form 23 January 2006; accepted 25 January 2006

Abstract

Graphical mechanisms for spatially organizing personal media data could enable users to fruitfully apply their conceptual models. This paper introduces *Semantic regions*, an innovative way for users to construct display representations of their conceptual models by drawing regions on 2D space and specifying the semantics for each region. Then users can apply personal categorizations to personal media data using the *fling-and-flock* metaphor. This allows personal media to be dragged to the spatially organized display and automatically grouped according to time, geography, family trees, groups of friends, or other spatially organized display representations of conceptual models. The prototype implementation for *semantic regions*, MediaFinder, was refined based on two small usability tests for usage and construction of user-defined conceptual models.

© 2006 Elsevier Ltd. All rights reserved.

Keywords: User interfaces; Personal media management; Spatial information management; Fling-and-flock; Dynamic queries

1. Introduction

Computer users often deal with large numbers of personal media objects such as photos, music, voice mail, video clips, web pages, emails, and document files. The term “personal media” suggests information sources that are familiar to users (e.g. family photos) and that have a modest number of items (say 100–10,000) so that browsing rather than search is

*Corresponding author. Tel.: +1 301 405 2725.

E-mail addresses: kang@cs.umd.edu (H. Kang), ben@cs.umd.edu (B. Shneiderman).

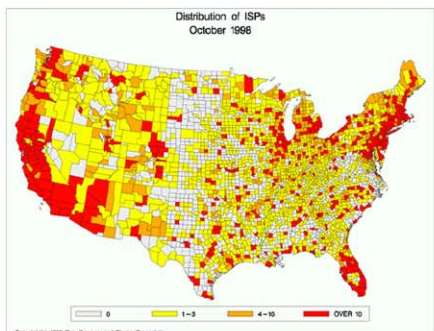
common. The process of browsing (“exploration of a body of information, based on the organization of the collections or scanning lists, rather than by direct searching” [1]) implies visual presentations that support exploration and promote retention. Managing personal media objects is a challenge for most users, who may struggle to understand, interpret, arrange, and use personal media objects. They wrestle with at least three major problems; an ever increasing amount of personal media data, rigid organizing metaphors, and difficulty in rapid data access.

Users understand their personal media by constructing conceptual models of it in their minds. There is no unique or right model. Rather, the models are personal, have meaning for the individual who creates them, and are tied to specific tasks. Even in a simple personal photo library, images can be organized by timelines, locations, events, people, etc. depending on users’ conceptual models. Despite the diversity of users’ conceptual models, the means available for users to organize and customize their information spaces are extremely poor and driven mostly by storage and distribution models, not by users’ needs. Consequently, one of the main challenges in designing a novel user interface for exploring personal media is to make it easier for end-users to construct, represent and apply their conceptual models flexibly for current tasks.


Another challenge in designing a personal media management system is to find ways to enable users to locate personal media objects rapidly. The wide and deep structure of file folder hierarchies as well as the inconsistent and sometimes meaningless file names often makes it difficult for users to find what they are looking for. Human capabilities for spatial cognition [2] can be a solution because they allow rapid information retrieval in 2D graphical user interfaces. Several experiments [3,4] have shown that spatial organizations of information enable users to access data items surprisingly quickly. Furthermore, if the spatial organization of information is based on familiar display representations or users’ conceptual models, task performance may improve even more.

Familiar display representations that facilitate comprehension include calendars, timelines, maps, organization charts, floor plans, body diagrams, building layouts, integrated circuit block diagrams, and periodic tables (Fig. 1). Unlike the standard mathematical or statistical visualizations such as histograms, line graphs, pie charts, scatter plots, lists, tables, and grids, most of these representations contain rich semantic information in their visual patterns. For example, organizational charts and genealogy trees include semantic information on the hierarchical structures of groups as well as their relationships. Timelines and calendars represent temporal information concerning events as well as time units.

The semantic combination of display representations and personal media can provide multiple perspectives for different information tasks. For example, suppose that a huge number of digital photos are combined with various display representations. The photos can be combined with a map display to show the distribution of photos according to the place they were taken. They can also be combined with a calendar display to represent the distribution of photos over days, weeks, months, or years depending on the time unit of the calendar. If the photos are about the anatomy of the human body, they can be combined with a human body diagram to specify which picture comes from which part of the human body. If the photos are about paintings taken from a museum, they can be combined with the museum layout diagram to indicate which painting was displayed in which exhibition room, or they can be combined with a timeline of art history to identify the trends of paintings at the time they were created. Likewise, other types of documents as




(a)




(b)



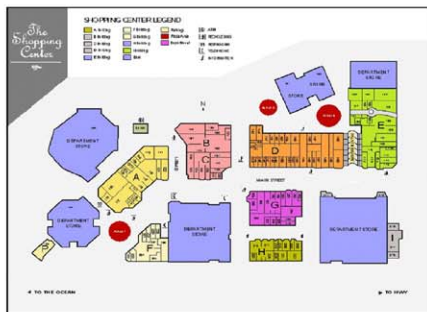
(c)



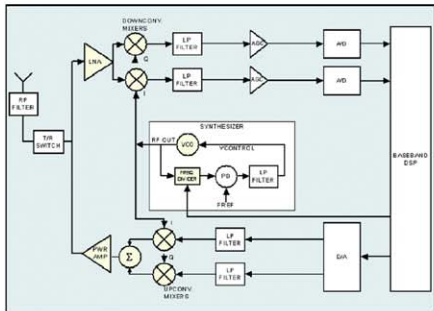
(d)



(e)



(f)



(g)



(h)

Fig. 1. Display representations of well-established and familiar conceptual models: (a) map, (b) floor plan, (c) organization chart, (d) body diagram, (e) periodic table, (f) building layout, (g) integrated circuit block diagram and (h) calendar.

well as personal information can be combined with visually compelling display representations to provide engaging perspectives for different information tasks.

These examples are well-established and familiar models that are shared by many people, but a successful tool should also enable users to create their conceptual models based on their perceptions of social networks, neighborhoods they live in, experiences in their lives, or their music library.

2. Related work

The large number of projects on personal media can be classified into three research directions: Personal Information Storage, Information Organization and Retrieval, and Flexible and Generalized Information Visualization.

2.1. *Personal information storage*

With the extraordinary progress made in digitization technologies [42], ranging from processor speed, storage capacity, data transmission rate, to scanner speed, camera resolution, audio/video encoding, display quality, and so on, there are many research efforts devoted to Personal Information Storage.

The Factoid project [5] imagined small ubiquitous devices that would capture many small facts per day, each only a few hundred bytes in size. Based on 1000 facts a day, they calculated 73 MB/year which can be saved forever and constitute a history of the users' life. In CyberAll, Bell [6,39] tries to archive all of his personal and professional information content including that which has been computer generated, scanned and recognized, and recorded on videotapes. He estimates that users will be able to record all of the information accumulated in their entire personal and professional life in a few terabytes. He focuses on issues surrounding this project such as longevity and long-term retrievability, access and access control, database and retrieval tools for non-textual information, and usability [40].

Other projects focused on augmenting memory. Pepys [7] used automatically captured location information from active badges to yield descriptions of past activities. It then combined these data into daily diaries of a user's activities. The main idea was that by providing these diaries, retrieval of information about the activities occurring over the course of the day would be improved. Video Diary enhanced the diary information presented by Pepys by providing automatically captured video images of the activities of the users. In Forget-Me-Not [8], information was automatically collected about a user's activities such as telephone calls made, e-mail notes received, and meetings attended. This information was stored and presented as a personal biography with activities presented in time order. The biography could be browsed, and filtered to show only a particular type of activity.

2.2. *Information organization and retrieval*

Technology progress and a growing desire by users to digitize and store their personal information have rapidly increased the amount of personal media data on individual machines. Therefore, indexing, retrieval, and management of personal media become critical issues. Several research attempts related to document management try to resolve similar problems by supporting a more flexible and intuitive organizational metaphor.

LifeStreams [9] uses a timeline as the major organizational metaphor for managing document spaces to replace conventional files and directories. LifeStreams is motivated by the problems of common desktop systems which have a standard single-inheritance file hierarchy. This approach creates difficulty in making use of archived information, getting an overview, and managing schedules and reminders. TimeScope [10] shows an interesting desktop metaphor called time machine computing (TMC) whose combination of spatial information arrangement and chronological navigation allows users to organize and archive information without being bothered by document folders or file classification problems. TimeStore [11] and LifeLines [12] are other time-based information indexing and retrieval systems. The designers of Stuff I've Seen [13] concluded that the multiplicity of independent applications used to manage information, each possessing its own organizational hierarchy (e.g., email, files, web, calendar), and the limited search capabilities often made it harder to find information on desktop, email store, or intranet than on the web. They tried to remedy this problem by providing a unified index across information sources. They explored new ranking and presentation ideas based on the assumption that rich contextual cues such as time, author, thumbnails, and previews can be useful because the information is personal and has been seen before. Later, they enhanced the Stuff I've Seen interface to represent the results of searches with an overview-plus-detail timeline visualization and extended a basic time view by adding public and personal landmarks.

These research projects attempted to solve problems of a standard single-inheritance file hierarchy by supporting a more flexible and intuitive organizational metaphor based on timelines. Although their organizing metaphor looks more intuitive and user-oriented, they simply replaced one superordinate aspect of the document (its location in the hierarchy) with another (its location in the timeline). Their organizing metaphors are still fixed and cannot be applied to problems that require multiple perspectives.

Several research projects used novel organizing metaphors, such as the Pile to support the casual organization of documents based on user perceptions of content [41]. Data Mountain [3] continued in this direction by allowing users to place documents at arbitrary positions on an inclined plane in a 2.5D desktop virtual environment using a simple 2D interaction technique. The benefits of human spatial memory were demonstrated in their study, but placement of documents on the Data Mountain was a tedious manual process.

Spatial layouts for conceptual models are widely used, often under terms such as Concept Maps [14]. The advocates suggest that tree and network layouts with connecting lines between nodes can represent complex issues in an understandable way. Concept Maps can be helpful as static portrait of a topic, but they do not have the interactive capabilities (e.g. region attributes and Boolean queries) that would organize personal media objects.

Tinderbox (<http://www.eastgate.com/Tinderbox>), Niagara [15], Visual Knowledge Builder [16] are visual information workspaces that support the collection, organization, and annotation of information. These tools make use of spatial cues as well as patterns in visual space to represent the characteristics of and interrelationships within the information. ContactMap [17] uses the models of personal social networks to provide an organizing principle for advanced user interfaces that afford information and communication services in a single integrated system. Since users invest considerable effort in maintaining links with networks of colleagues, acquaintances and friends, the display presentation of social networks is valuable. Although spatial organization of

information enabled users to retrieve documents quickly, the ContactMap did not allow users to change the organizations based on their tasks or conceptual models, nor represent a known organizational hierarchy.

A few research efforts have indexed personal media with high-level attributes to construct flexible hierarchies or multiple categorizations which provide users multiple viewpoints for personal information. Presto [18] provides rich interaction with documents through user-level document attributes. Document attributes capture the multiple roles that a single document might play, and allow users to reorganize their document space for the task at hand. The model enables users to interact with their documents using high-level attributes that are fluidly managed by the users themselves. The Semantic File System [19] provides flexible associative access to the system's contents by automatically extracting attributes from files with file-type specific transducers. Virtual directory names are interpreted as queries and thus provide flexible associative access to files and directories in a manner compatible with existing file systems. Haystack [20,21] uses a semistructured data model to describe the connection between different documents in a user's corpus as well as the metadata concerning each document. It adopts a category-based organization and navigation scheme that allows information to be placed in multiple thematic bins or categories, simultaneously. This research is based on the belief that information inherently has multiple, relevant categories that users can readily identify and that allowing multiple categories lets users organize documents in a more intuitive, richer information space.

2.3. Flexible and generalized information visualization

Users often need multiple views to handle the different types of information and multiple perspectives in a complex information task [22,17]. However, it is difficult for users to find generalized visualization tools satisfying their demands, or to change visualizations dynamically according to their conceptual models or specific tasks.

Many generalized information visualization and visual data-mining tools are available, both commercial and experimental, including Spotfire (www.spotfire.com). Popular features of Spotfire include direct data manipulation, dynamic queries, easy import/export of data, rapid change to axes, and flexible color/size coding [23–25]. However, it has a limitation in that all data must be projected onto the mathematical 1–3D spaces as bar charts, pie charts, or scattergrams even though they can be visualized more intuitively in other simple user-friendly display representations such as calendars, genealogy trees, organizational charts, etc. Like other major statistical data mining tools such as SAS JMP, DataDesk, and SPSS Diamond, Spotfire is more appropriate for visualizing scientific and quantitative data rather than metadata-based personal media data [26].

Visage [27] is a prototype user-interface environment that enables users to move and combine interesting subsets of information across the isolating boundaries imposed by different applications. Visage and SageBrush [28] enables users to assemble combinations of graphical elements (e.g. axes, lines, text, etc.) and map these elements to the data that users wish to visualize. Snap-together [29] visualization enables users to rapidly and dynamically construct coordinated-visualization interfaces, customized for their data, without programming. Users load data into desired visualizations, then create coordination between them, such as brushing and linking, overview and detail, and drill down. However, they have a limitation that their data should be well-structured and centralized, which is not appropriate for personal media data.

On the other hand, there was an early paradigm of separating data and visualization to enable more flexible information visualization. In the Model/View/Controller (MVC) [30] paradigm, the user input, the modeling of the external world, and the visual feedback are explicitly separated and handled by three types of object, each specialized for its task. The “view” manages the graphical and/or textual output to the portion of the bitmapped display that is allocated to its application. The “controller” interprets the mouse and keyboard inputs from the user, commanding the model and/or the view to change as is appropriate. Finally, the “model” manages the behavior and data of the application domain, responds to requests for information about its state, and responds to instructions to change its state. The formal separation of these three tasks makes GUI design and implementation more flexible and powerful and also shows the possibility of applying various views to the same model according to the users’ conceptual models and personal categorizations.

3. A spatial user interface supporting user-defined conceptual models

The rich sets of research projects described in Section 2 demonstrate the strong interest in solving personal media management problems. The contribution in this paper is a novel combination of tools that enable users to create flexible display representations as screen regions and automatically place personal media items in appropriate regions. Our *semantic regions*, stems from one central assumption: Spatially organized information based on the semantics of personal media and the users’ conceptual models for managing personal media will improve task performance as well as user satisfaction. We implemented the MediaFinder [31] prototype to investigate the use of *semantic regions* for managing and exploring personal media.

3.1. The basic concept

Computer users usually understand personal media data by constructing their own conceptual models. For example, with music, people use different classifications depending on their interests and tastes. Some people may classify music into three general categories such as pop, classical, and folk, but others may break it down into more detailed categories such as rock, jazz, rhythm and blues, rap, country, classical, latin, and so on. Even in the classical category, some may conceive of music by historical periods such as Renaissance, Baroque, Classical, Romantic, and Modern, while others may classify music into genres such as ballads, canons, concertos, etudes, fantasies, fugues, inventions, oratorio, requiems, sonatas, symphonies, and so on.

In addition, although they have similar categorizations of music, their conceptual maps can be differently organized and explicitly represented on a 2D space. For example, some people may see rhythm and blues as closer to rock, whereas others see rhythm and blues as closer to country and to jazz. Furthermore, they may make use of size and color to show the strength of their interests or number of items they have (e.g. rock might be twice as large as rhythm and blues, and four times as large as jazz). They may even make use of shapes or connecting lines to make their conceptual maps more understandable.

A second example might be the personal cognitive map people have of their neighborhood. Unlike an actual map, the personal cognitive map shows strength of interest in a local park, a neighbor’s house, a grocery store, a swimming pool, a library, etc.

Although their maps may be distorted and have a few different scales depending on their strength of interest or the amount of information related to the places in the map, it helps them understand, recognize and explore information in their own way.

Because of the diversity of people's conceptual models, it is not easy for users to manage their personal media data with a fixed organizing metaphor as provided in current file systems. The *semantic regions* approach provides users with an environment to customize their information space appropriately for their conceptual models and specific tasks.

3.2. Semantic regions

Semantic regions are query regions drawn directly on a 2D information space. Users can specify the shapes, sizes, and positions of the regions and thus form the layout of the regions meaningful to them. *Semantic regions* created with MediaFinder (see examples in Fig. 2) are spatially positioned and grouped on the 2D space based on personally defined clusters or well-known display representations such as a map, tree, timeline, or organization chart. *Semantic regions* can be hierarchically organized, such as in the world map that has country regions, with state regions, and a state region that has city regions (Fig. 2(a)). Once the regions are created and arranged, users can specify semantics for each




Fig. 2. Example models of *semantic regions* constructed with MediaFinder: (a) world map model, (b) genealogy tree model, (c) season calendar model and (d) group organization chart model.

region. The semantics can be composed of file attributes, user-defined high-level attributes, or database attribute depending on the types of personal media data. The attributes are conjunctively joined to form the semantics of a region.

In Fig. 3, each region represents a person, and the regions are grouped into 5 clusters to represent different friend groups. When photos are dragged onto the main panel, they are automatically placed in the appropriate regions based on the annotations.

This metaphor is called *fling-and-flock*; that is, users fling the objects and the objects flock to the regions. If photos do not satisfy any of the semantics of regions, they are collected in the *remaining items region* located at the top of the panel. Users can save the current state of regions and subsequently apply it to other personal media data sets.

A semantic region has some similarities with a folder, since multiple personal media items are contained in a semantic region just as multiple files are contained in a folder. *Semantic regions* can have hierarchies. Child regions can be defined inside a parent region like subfolders. However, there are two major differences between a semantic region and a folder. First, a semantic region has explicit semantics and it only contains the personal media items that satisfy the semantics. A semantic region automatically gathers the personal media items based on its semantics, and users cannot enforce the personal media

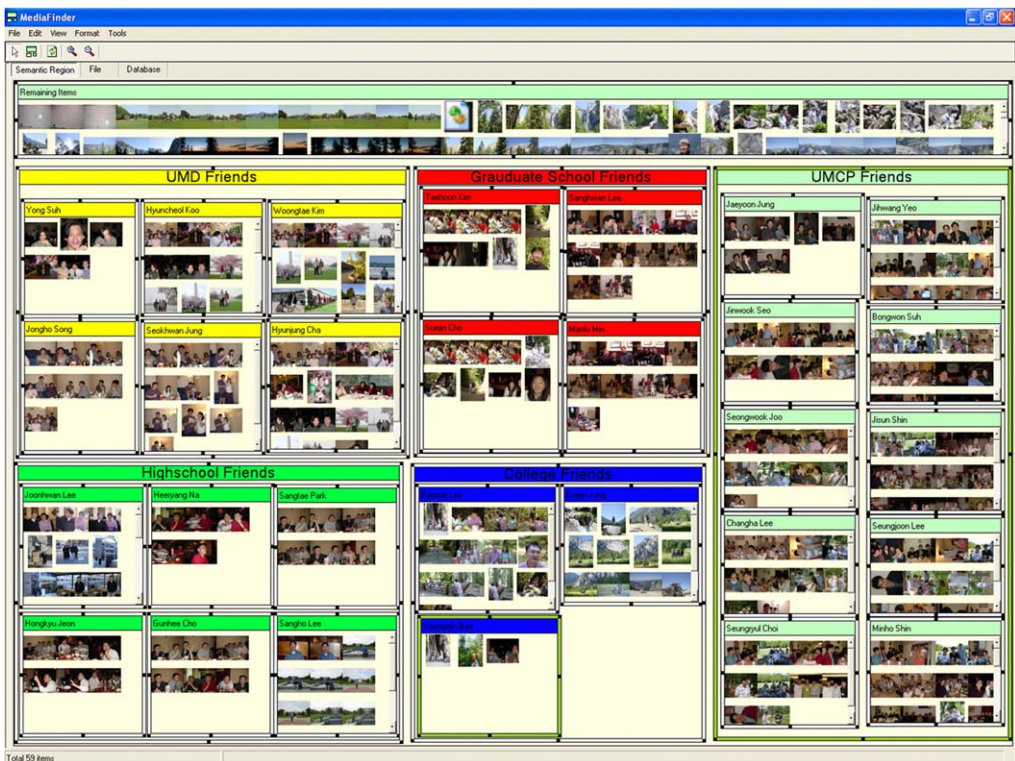


Fig. 3. A friend group conceptual model: each region represents a person and contains all the photos annotated with the name defined in it. The regions are grouped into 5 clusters to represent different friend groups (UMD friends—University of Maryland computer science students, high school friends, graduate school friends, college friends, and UMCP friends—University of Maryland non-computer science students). Each group has its own color to represent the different group of friends.

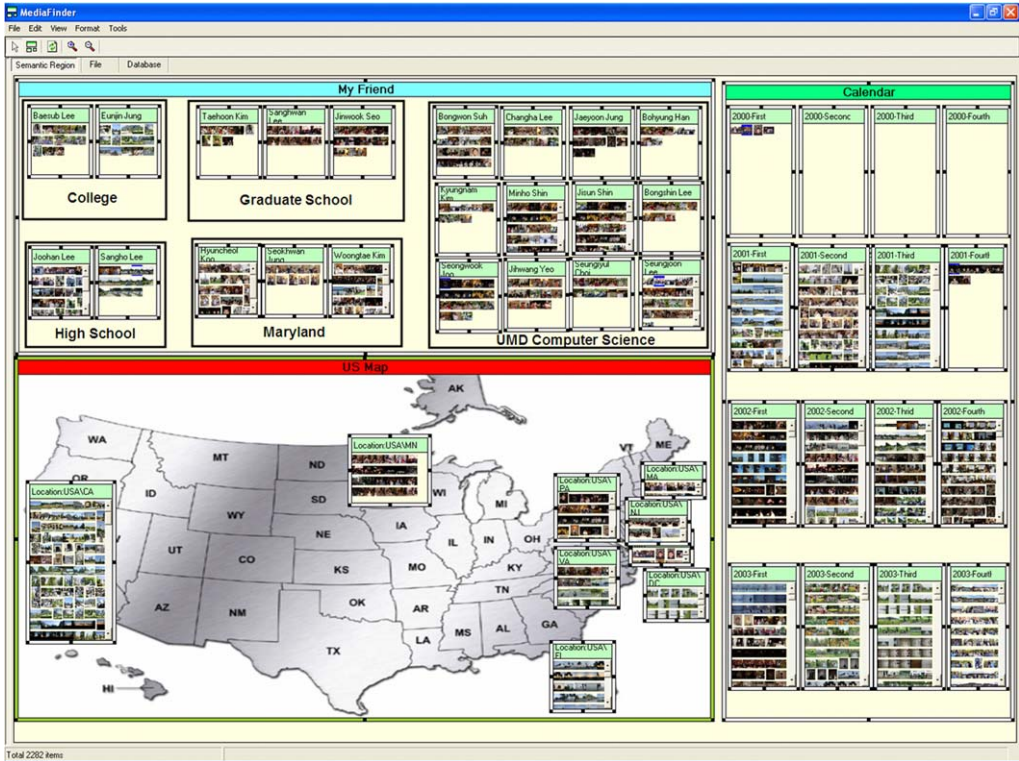


Fig. 4. Three conceptual models, friend groups, quarter year calendar, and US map (clockwise from the top region) are combined together. The regions in the My Friend region are grouped into 5 clusters to represent different friends groups (high school friends, college friends, graduate school friends, UMD CS friends, and MD friends). Each region in the calendar group represents a quarter year from 2000 to 2003.

data to be placed in a region unless the attributes of personal media items are changed. Second, a personal media item can be contained in multiple regions as long as it satisfies their semantics, thus enabling more flexible organizations.

Using *semantic regions*, users can apply several conceptual models simultaneously. For example, in Fig. 4, three conceptual models, the US map, the year calendar, and my friends group are combined, all showing the same photos.

3.3. MediaFinder

The *MediaFinder* is a prototype interactive tool built to investigate the use of *semantic regions* for personal media management and exploration. MediaFinder provides a working environment for *semantic regions* construction and operation. Also, MediaFinder functions as a repository to hold a domain personal media data set for data binding. Users can bind personal media data dynamically through the *fling-and-flock* metaphor. In addition, MediaFinder provides users a working environment to represent their conceptual models, combine them, and use them for various personal media management tasks such as organization, meaning extraction, search, navigation, indexing, and distribution.

3.4. Construction of semantic regions

MediaFinder is composed of a 2D workspace and functions that enable users to construct and use *semantic regions*. In the workspace, users can freely create, remove, resize, move, and change the z-order of *semantic regions* just as in object-oriented drawing tools. After users draw the regions to represent their conceptual models, they can specify the semantics for each region.

3.4.1. Semantic region creation

The creation of *semantic regions* is straightforward: users click on the desired starting point of the region and drag the pointer to the desired location of the opposite corner. Once the region is created, it may be dragged to a new position or resized via appropriate resize handles on the corners. Users can change the shape of the selected regions, connect them with lines, or add background images to clearly represent their implications and thus easily remember and recognize the constructed model. Child regions can be constructed recursively inside a region. The semantics of the child regions are applied only to the personal media items that are contained in their parent region. Hence, if the semantics of the parent region are modified, then the personal media items contained in the child regions are dynamically changed.

Each region has multiple views to provide alternative perspectives on the contained personal media item. Each region supports five basic views (list, small icon, large icon, details, and thumbnails) and three additional views (aggregation, intersection, and background) to provide an overview and represent the relationships among regions. In the aggregation view, the region shows the number of personal media items that it contains. This information can be encoded as a color or a size of the region to assist quick perception. The background view enables users to easily remember the implication of the region through a background image and an overview of the child regions. The intersection view represents the number of personal media items shared among regions.

The overall design of multiple views for *semantic regions* was based on the well-known information-seeking mantra, “Overview first, zoom-in, then filter, and detail-on-demand” [22]. Users can see an overview of the distribution of personal media data first, and then see more details by zooming into a region. The design was also based on query previews [32] and dynamic aggregation [33–35].

Users can draw lines connecting the semantic regions to illustrate the interrelationship among the regions more clearly. Users can group/ungroup the semantic regions to create/remove a hierarchy among the regions. MediaFinder provides several align and format functions to facilitate the construction of users’ conceptual models.

3.4.2. Specification of semantics

Each semantic region has semantics to describe the personal media items. In *semantic regions*, we defined the term “Semantics” as any metadata that the personal media can have. The “Semantics” can be the file attributes of MS Windows file system (e.g. file name, path, date created/modified, size, type, etc), the user-defined high-level attributes (e.g. location, rating, people, season, event, etc.), or the attributes of a database table. The basic assumption in designing MediaFinder is that personal media objects already have various kinds of pre-defined metadata that can be used for specifying the semantics of a semantic region. Our assumption is based on the fact that the current technologies allow users to

extract metadata from their personal media objects in various ways; such as automatic online creation (e.g. the attribute values of emails and files, time stamp or GPS tag of digital photos, etc.), automatic offline extraction (e.g. face recognition, human detection, keyword extraction, automatic video indexing, etc.), automatic transfer from other sources (e.g. Internet databases for movies and CDs) and also including manual creation using the advanced annotation tools (e.g. PhotoFinder [36]). The technology progress will accelerate the extraction as well as the creation of metadata so that users can have ample metadata to be used for personal media management. Based on this assumption, MediaFinder has been designed to make use of personal media metadata to form the semantics of a region. In this section, the personal digital photos annotated with PhotoFinder are used as personal media objects to simplify the example of specifying semantics to the regions.

Right-clicking on a region pops up a menu, from which users can bring up the *semantic regions attribute dialog box* to specify the region name as well as the semantics to the region. In this example (Fig. 5(a)), the metadata of personal photos are classified into four tabs; File, Date, People and Category. Each tab shows the list of metadata collected from the personal photos so that they can be selected and joined together conjunctively to specify the semantics of the region. In the semantic regions attributes dialog box, File and Date tabs represent the file attributes of personal photos, which were automatically extracted from the file system. On the other hand, Category and People tabs represent the user-defined attributes that were manually annotated with PhotoFinder. In the category tab, selecting a node from the category tree (left window) and adding it to the selected category listbox (right window) specifies the semantics (Fig. 5(a)). Although the “People” can be included in the user-defined “Category”, MediaFinder was designed to have a separate “People” tab (Fig. 5(b)) because the names of people are frequently used as key metadata in personal photo management.

All the categories and people attributes specified for a region are combined conjunctively. In other words, only the personal photos that have all of the specified attributes can be contained in the region. For example, if a region in MediaFinder is

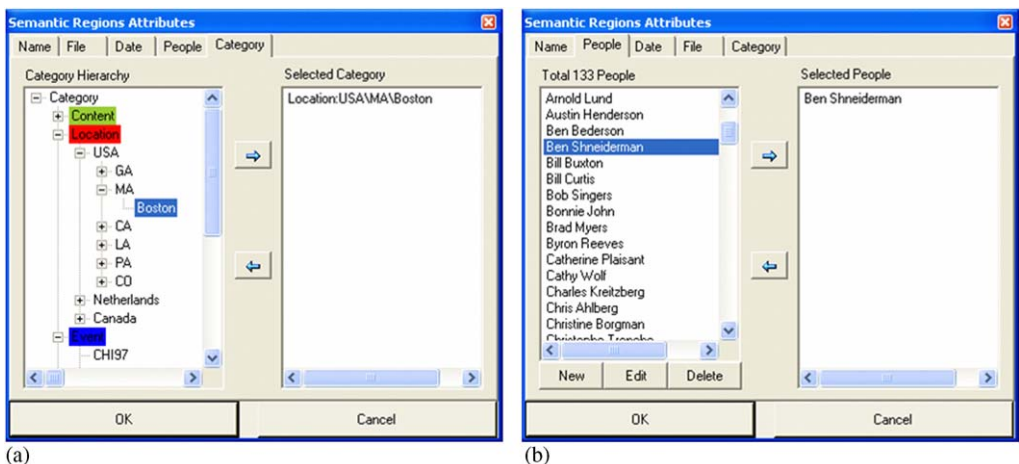


Fig. 5. Specifying the user-defined attributes for a region in the category and the people tab of the semantic regions attribute dialog box: (a) specification of category attributes and (b) specification of people attributes.

specified with the category attribute “Location:USA\MD\College Park” and the people attribute “Hyunmo Kang”, then only the photos of “Hyunmo Kang” taken in “College Park, Maryland” will be contained in the region and other photos will be filtered out from the region.

MediaFinder makes use of Microsoft Windows NTFS (Windows NT File system) file summary property to keep the user-defined attribute values of personal media items. NTFS provides the file summary properties so that users can read and write a description of a file and these properties are then used to store and retrieve the user-defined attribute values in MediaFinder. When personal media items are bound with MediaFinder through *fling-and-flock*, MediaFinder collects all the user-defined attribute values stored in NTFS file summary properties of the personal media items and builds the hierarchies of category and people attributes (see Fig. 5) to represent in the semantic regions attributes dialog box.

3.4.3. Model management

The layout of the regions as well as the semantics of each region can be saved as a file for later use. Users can apply multiple conceptual models to the same personal media data set by opening the previously saved files. Also, users can import and combine multiple models in MediaFinder to discover complex relationships among items and categories.

The models can be combined horizontally or vertically (Fig. 6). Users can combine the multiple models on the same level so that all the models are applied to the same personal media items simultaneously. By combining the models horizontally, users can easily recognize the relationships among the models and formulate a query across the different models. On the other hand, users can import a model inside a region so that the imported model is applied only to the personal media items that are contained in the parent region. By combining the models vertically, users can dynamically regroup the personal media items in a region based on a variety of models.

The MediaFinder prototype was designed to have a four-level hierarchy overall: Model–Semantic Region–Personal Media Item–Metadata. In other words, a model is composed of several semantic regions where each region contains personal media items that are annotated with the metadata.

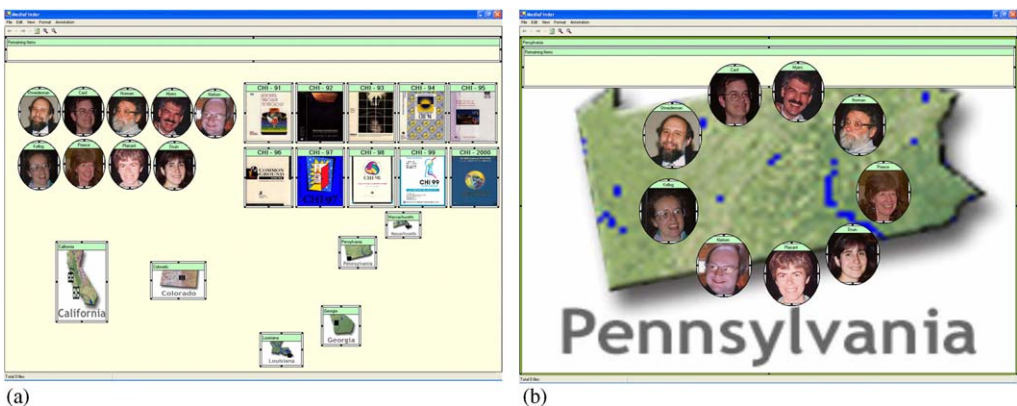


Fig. 6. Multiple models can be combined horizontally or vertically: (a) three models combined horizontally and (b) people model inside a location region.

3.5. Using semantic regions in MediaFinder

Once the *semantic regions* are defined, MediaFinder users can organize their personal media items and explore them.

3.5.1. Data binding (the *fling-and-flock* metaphor)

Once the *semantic regions* are created, they are ready to be applied to personal media items. When users drag and drop the personal media items from the Windows file explorer onto the model, the items are automatically distributed and placed in the appropriate regions based on the semantics defined in each region. If the personal media items do not satisfy any of the semantics of the regions, they are collected into the remaining items region (Fig. 7).

During *fling-and-flock*, MediaFinder collects all the attribute values of the bound personal media items. The collected attribute values are used for specifying the semantics for a region or for giving an overview of what kinds of metadata are annotated to the personal media items just as query preview [32] does. The collected user-defined attributes are translated to form a semantic hierarchy as shown in Fig. 8. It shows that a total of 6 user-defined categories were collected from the personal media items and each category had its own hierarchy. For example, the location category has a three-level hierarchy

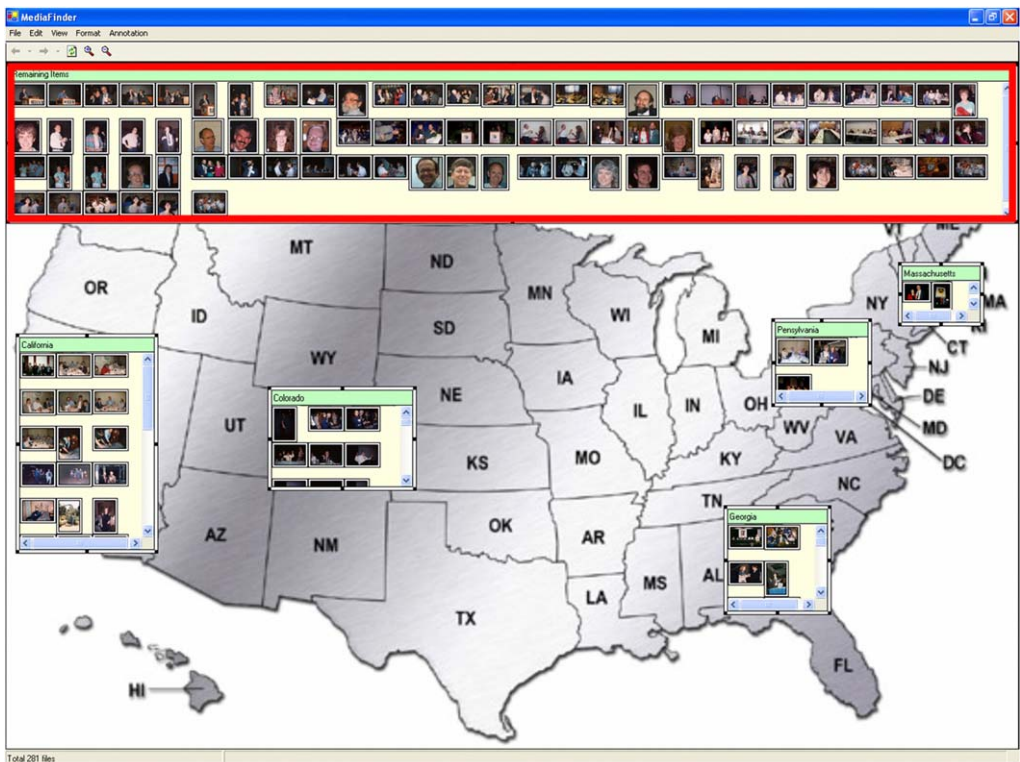


Fig. 7. The photos that do not satisfy the semantics of five state regions are collected in the *Remaining Items* region (red rectangle) located at the top of MediaFinder main window.

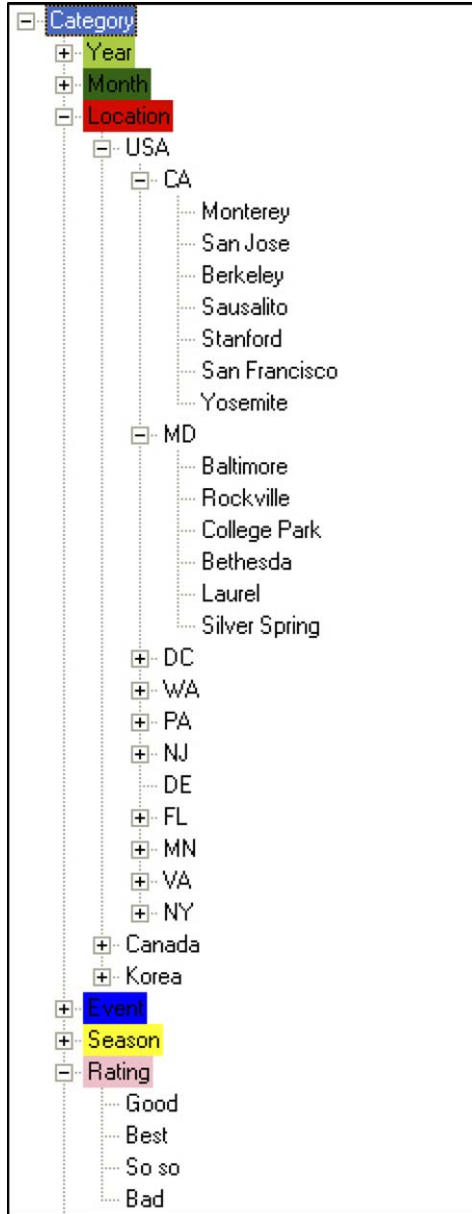


Fig. 8. User-defined attributes collected by MediaFinder through fling-and-flock.

(country–state–city), but the Rating category has just one-level hierarchy to represent the rating of the personal media items.

3.5.2. Flexible organization

Each personal media item can be contained in multiple semantic regions as long as it satisfies the criteria of the regions. The semantics of the regions do not need to always be

mutually exclusive. This design is based on the fact that users’ conceptual models for categorizing personal media are not always mutually exclusive unlike the file system structure. For example, suppose that users want to organize their digital photos based on people who appeared in the photos. Then, the photos in which more than one person appeared should be contained in multiple collections. This flexible organization of personal media items also makes it possible for users to combine different conceptual models and observe the relationships among them.

3.5.3. Navigation

Semantic regions can be zoomed in to show the detail information of the contained personal media items (Fig. 9). Fig. 9(a) shows the USA map model that contains five state regions. A region is zoomed in to fill the parent region with animation when it is double-clicked. Fig. 9(b) shows the zoomed in California region with a detail view to display the list of the files contained in the region. The background view shows the child regions located inside the California region.

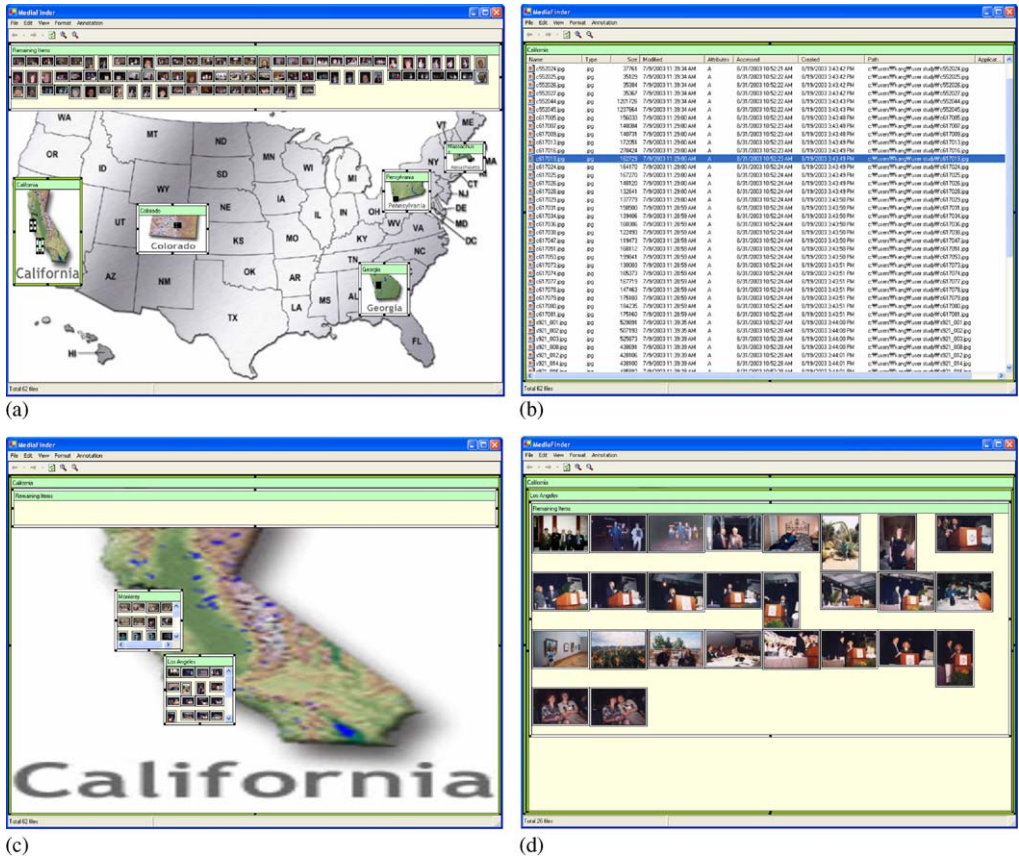


Fig. 9. The zoomed in region shows more details about the contained personal media items or represents the hierarchy of child regions: (a) zoom in California by double-click, (b) detail content list in California region, (c) two city child regions in California and (d) zoom in a city region recursively.

In Fig. 9(c), the personal media items contained in the California region are redistributed and placed in the two city regions (Los Angeles and Monterey). The Los Angeles city region is zoomed in recursively and shows the contained items with a thumbnail view (Fig. 9(d)). The titles of the entire zoomed in regions are displayed contiguously to represent the current level in the hierarchy (Fig. 9(d)). The zoomed in region shrinks back to its original size when it is double-clicked again. All of the zooming processes are animated to reduce the users' confusion that might be caused by the change of region hierarchy.

3.5.4. Meaning extraction

MediaFinder provides an interactive visualization technique called *region brushing* to visualize the relationships among the semantic regions. Region brushing is used to highlight the personal media items contained in multiple regions simultaneously and is often used for depicting both the intra-relationships and the inter-relationships among models. Since a personal media item can be contained in multiple semantic regions, the existence as well as the amount of the shared personal media items among the regions well represents their relationships. Fig. 10(a) shows a “HCI Researchers” model in which each region represents a person. The conference photos were bound to this model through *fling-and-flock* so that each photo could be placed in multiple people regions based on its annotations. Region brushing works as follows. When the mouse is over an item in a region, borders of the regions that contain the item are highlighted red. The selected item in the highlighted regions is also highlighted so that it is easily identified (Fig. 10(b)). On the other hand, when the mouse is over the title of a region, all the regions that contain any of the items in this region are highlighted red while highlighting the containing items (Fig. 10(c)).

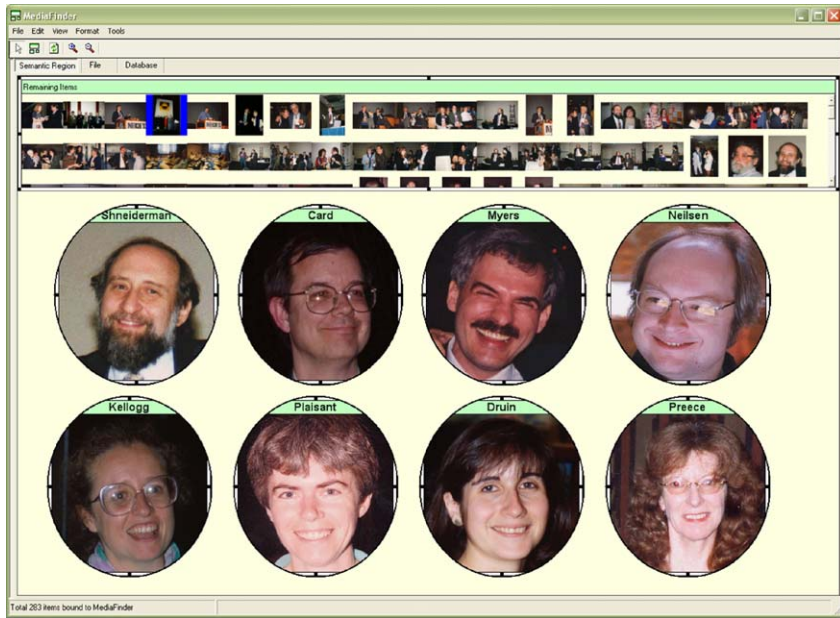
Semantic regions were designed to provide an “Intersection view” to visualize how many personal media items in the selected region are shared with the other regions (Fig. 10(d)). With this interactive visualization technique, users can quickly recognize the relationships among regions and thus answer questions like, “Who appeared in pictures with Ben Shneiderman?”, “How many pictures were taken with Ben Shneiderman?”, or “What percentage of the pictures of a person were taken with Ben Shneiderman?”.

The region brushing technique becomes even more powerful when multiple models are combined horizontally to represent interrelationships among the models in MediaFinder. In Fig. 11, three models, HCI researcher, US map, and CHI conference calendar are combined horizontally and bound to the conference photos through the *fling-and-flock* metaphor. Since a photo can be contained in multiple regions across the different models, many questions concerning the interrelationships of models can be answered through region brushing. Such questions might be “Find the name of the conferences that all nine people participated in and where was the conference held?”, “What was the name of the conference held in Atlanta, Georgia, and who did not appear in the photos taken at this conference?”, “Find the conferences in which Ben Shneiderman did not appear in the photos and where was the conference held?”, “Find the conferences that were not held in the US”, and so on.


3.5.5. Search

In MediaFinder, users can apply Boolean operations (AND, OR, and NOT) or set operations (Intersection, Union, and Complement) to *semantic regions*. When an AND operation is applied to the selected regions, a new region is created over the regions to

contain only the personal media items that are contained in every selected region. Likewise, if the OR operation is applied to multiple regions, a new OR region is created in order to contain the personal media items that are contained in any of the selected regions.




(a)

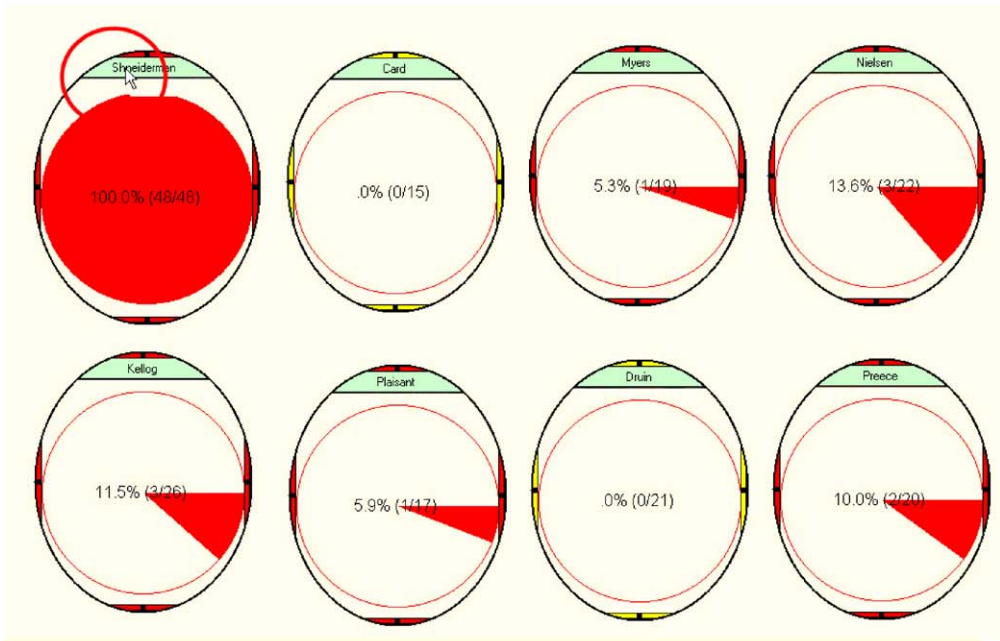


(b)

Fig. 10. Region brushing: an interaction technique to visualize the shared items among the *semantic regions*: (a) “HCI researchers” model bound to the conference photos, (b) the selected item and all the containing regions are highlighted red, (c) regions containing any of the items in the mouse-overed region are highlighted and (d) intersection view to represent the number of shared items among the regions.



(c)



(d)

Fig. 10. (Continued)

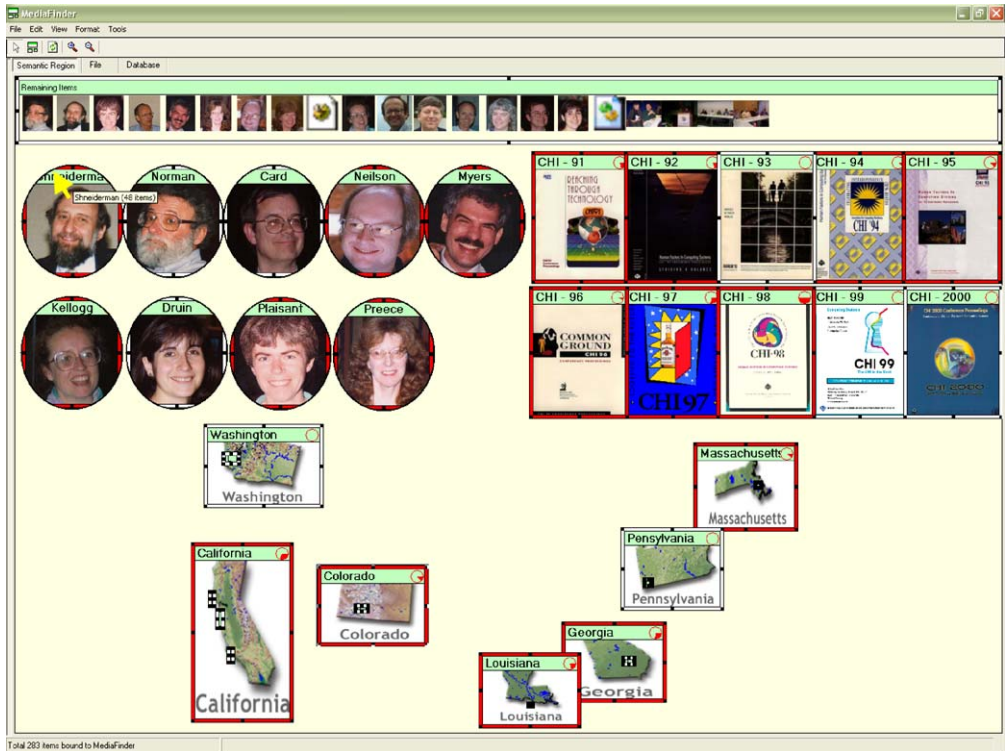


Fig. 11. Extracting meanings concerning the interrelationships of the multiple models through the region brushing technique. By placing the mouse on Shneiderman's region, users can see which conference and which state his photos were taken.

If a NOT operation is applied to a region, a new region is created right over the selected region to contain all the personal media items that are not contained in the selected region but bound to MediaFinder. The Boolean operations can be recursively applied to all regions. With the Boolean operations, users can specify any logical combination of the region semantic.

Fig. 12 shows the process of applying Boolean operations to the regions. For example, consider the query “Find all the photos in which either Ben Shneiderman or Jenny Preece appeared but not with Catherine Plaisant”. A total of three Boolean operations are needed to achieve this goal. In Fig. 12(a), the procedure starts with applying the NOT operation to the region of Catherine Plaisant by choosing a “Complement” menu from the popup menus. Then the NOT region is created to contain the photos that are bound to MediaFinder but not annotated with the name, Catherine Plaisant (Fig. 12(b)). The third step is to select both the Ben Shneiderman region and the Jenny Preece region to apply an OR operation (Fig. 12(c)). The OR region is created over the selected regions and it includes all the photos that are annotated with either Ben Shneiderman or Jenny Preece (Fig. 12(d)). Finally, the AND operation is applied to the NOT and the OR regions (Fig. 12(e)), and the AND region is created to contain the photos that only satisfy the Boolean query (NOT (Catherine Plaisant)) AND ((Ben Shneiderman) OR (Jenny Preece)).

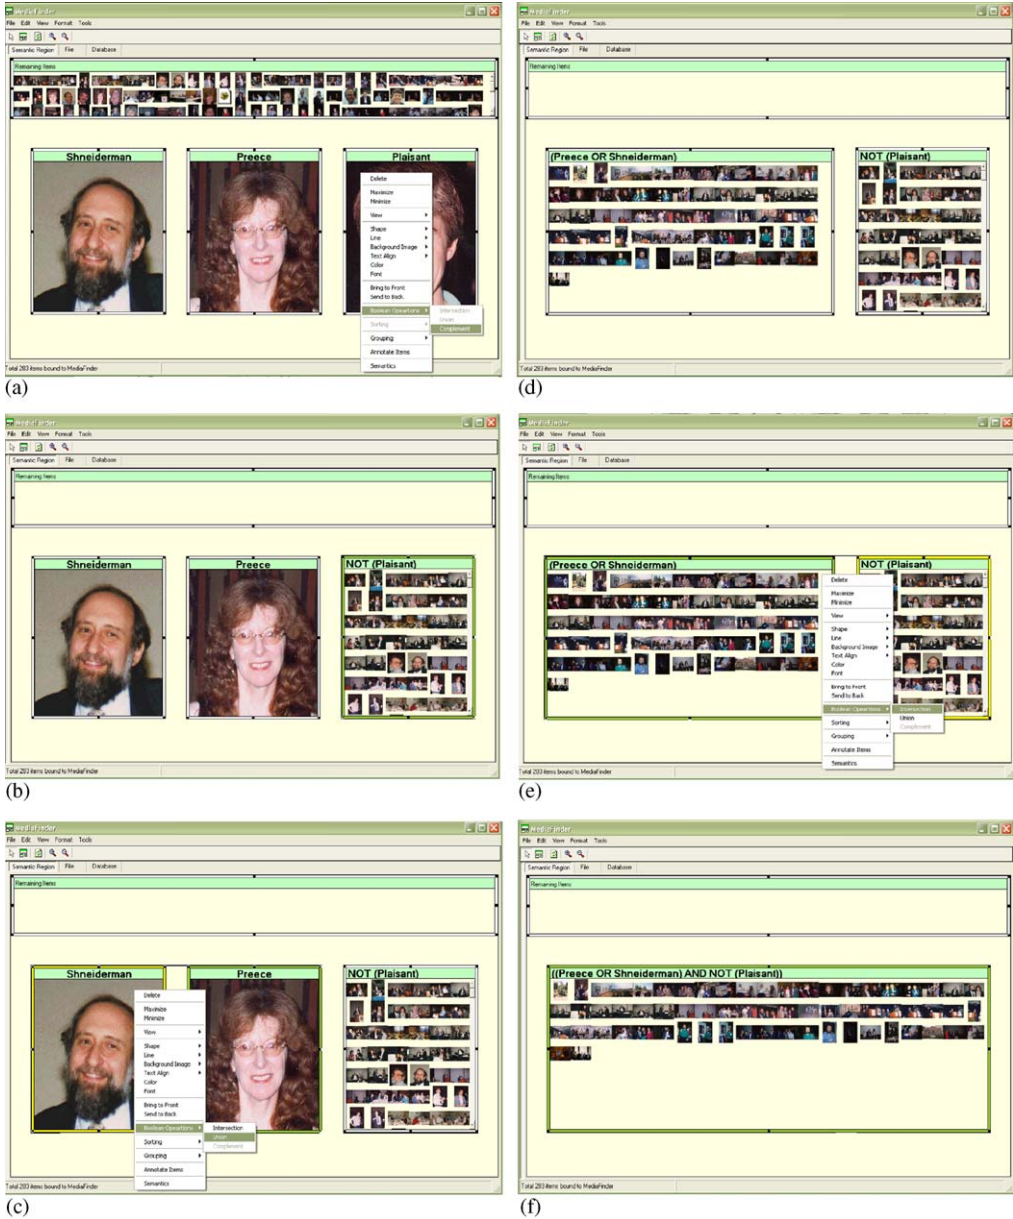


Fig. 12. The steps for applying Boolean operations to get the region of (NOT (Catherine Pleasant)) AND ((Ben Shneiderman) OR (Jenny Preece)): (a) *Step 1*: Apply NOT to the region of Catherine Pleasant, (b) *Step 2*: A new NOT region is created over the region, (c) *Step 3*: Applying OR to Ben Shneiderman and Jenny Preece regions, (d) *Step 4*: A new OR region is created over the regions, (e) *Step 5*: Apply AND to two regions and (f) *Step 6*: Final AND region constructed over the OR and the NOT regions.

These newly created regions do not have any user-specified semantics but they contain the information about what Boolean operation was applied and what regions were used as operands for this operation. If there is any change in the contents of an operand region

caused by modifying the semantics of the region or binding a new set of personal media items, all the regions that use this region as an operand are updated together immediately. Cycles are not allowed in the hierarchy of Boolean operations.

3.5.6. Indexing and distribution

Even though MediaFinder has been designed to use the pre-defined personal media metadata, *semantic regions* can also be used for indexing/annotating personal media items manually. If users select personal media items in one region and drag them to another region, the attribute values of the selected items are updated to those specified in the destination region. If the semantics of the destination region and the original region are not mutually exclusive (in other words, if a personal media item can be contained in both of the regions simultaneously without any conflict of the semantics), the personal media items are duplicated and contained in both regions. On the other hand, if the semantics in the destination region conflict with the semantics of the original region, the personal media items are removed from the original region and added to the destination region.

In MediaFinder, the drag-and-drop metaphor is used for indexing personal media items just as it is used in direct annotation [37]. In Fig. 13(a), five photos in the remaining items region are selected and dragged onto the Minnesota region. If the photos are dropped onto the region, they are annotated with the location attribute value “location:USA\Minnesota” that is specified in the Minnesota region. Subsequently, the dragged photos are removed from the remaining items region and added to the Minnesota region because they now satisfy the semantics of the Minnesota region (Fig. 13(b)).

Another major personal media management task is to distribute the personal media items either in private or in public. MediaFinder was designed to provide a few distribution features even though they were not the focus of this research. The basic unit of distribution in MediaFinder is a region. Users can copy the personal media items in a region to the computer file system by dragging them to a folder or through the copy-and-paste function. *Semantic regions* were also designed to be associated with the default email client or ftp client to export the items.

4. Evaluation and refinement

Studying the use of semantic regions is important for three reasons: To evaluate the usability and benefits of MediaFinder, to discover potential user-interface improvements, and to gain a deeper understanding about users’ ability to understand, construct, and use semantic regions. Two small usability studies were conducted: the first to determine whether users can understand the concept of semantic regions and the second to observe whether users can construct semantic regions based on their own conceptual models. The studies mainly focused on the qualitative user studies in terms of user satisfaction rather than task performance. This is mainly because there was no available personal media management tool to compare the task performance with MediaFinder other than file systems or database systems. Although we use a database system for task performance comparison, it is not reasonable to expect general users to make SQL queries that are not easy to compose even for the highly experienced computer users. For example, if the region



(a)



(b)

brushing task shown in Fig. 10(c) and search task in Fig. 12 can be translated to SQL queries as follows:

- Region Brushing Task in Fig. 10(c):

```
SELECT Photos.PhotoID, Photos.URL, Appearance.PersonID
FROM Photos INNER JOIN (People INNER JOIN Appearance ON People.PersonID =
Appearance.PersonID) ON Photos.PhotoID = Appearance.PhotoID
WHERE (((Photos.PhotoID) In (SELECT Photos.PhotoID
FROM Photos INNER JOIN (People INNER JOIN Appearance ON People.PersonID =
Appearance.PersonID) ON Photos.PhotoID = Appearance.PhotoID
WHERE (((People.FamilyName) = "Shneiderman"))
GROUP BY Photos.PhotoID)))
GROUP BY Photos.PhotoID, Photos.URL, Appearance.PersonID
ORDER BY Appearance.PersonID;
```

-Search Task in Fig. 12:

```
SELECT Photos.PhotoID
FROM Photos INNER JOIN (People INNER JOIN Appearance ON People.PersonID =
Appearance.PersonID) ON Photos.PhotoID = Appearance.PhotoID
WHERE (((((People.FamilyName) = "Shneiderman" Or (People.FamilyName) = "Preece")
And (Not (People.FamilyName) = "Plaisant"))))
GROUP BY Photos.PhotoID;
```

As shown in the above-complicated SQL query examples, it is reasonable to assert that the MediaFinder graphical user interface enables general computer users to perform the personal media management-related tasks much more easily and with less training than with SQL. Even for professional users, MediaFinder should improve productivity by speeding performance and reducing errors.

Fig. 13. Indexing the personal media items by dragging them onto the semantic region: (a) 5 photos are dragged from remaining, (b) dragged items are annotated with the items region onto the Minnesota region and semantics of the Minnesota region.

4.1. Evaluation of semantic regions use

Background information was obtained from the five subjects concerning their experience with computers. The 10-min training consisted of three phases: a quick demonstration of MediaFinder by the administrator to give an overview with motivation (5 min), review of personal media management concepts such as organization, meaning extraction, search, and navigation (2 min), and detailed instruction on the usage of MediaFinder and semantic regions (3 min).

Then, when confident to continue, each subject was given a set of practice tasks, and allowed to ask questions about MediaFinder. The test data set consisted of 281 photos from CHI conferences between 1991 and 2000. The photos were stored in one directory and their metadata (e.g. location, date, and people) was annotated with PhotoFinder. In addition to photos, four semantic regions templates were provided: Location Map (7 states in the US with 8 cities), Conference Calendar (10 years), People (9 most frequently appearing people), and Combined Model (combination of the above three models). Seven tasks were given to the users.

- Location Model

1. Group the photos based on the US map. Which state has the most photos and which state does not have any photos? How many photos were taken outside the US? (Organizing photos using the *fling-and-flock* metaphor.)

- Calendar Model

2. Regroup the photos based on the Conference Calendar model. How many photos were taken in each conference? (Dynamic regrouping of personal media: Apply a different conceptual model to the same data set.)

- People Model

3. How many people appeared in at least one photo along with Ben Shneiderman? (Understand the multiple existence of an item in the regions.)

4. Find all the photos where both Shneiderman and Kellogg appeared. (Search: apply Boolean operations to the regions.)

- Combined-model

5. Find the name of the conferences that all nine people participated in. Where (what is the name of the state) was the conference held? (Find out the relationship among the different conceptual models.)

6. What was the name of the conference held in Atlanta, Georgia? Find out who did not appear in the photos taken at this conference. (Change the focus among different models.)

7. Find the conferences that were not held in the US. (Extra meaning extraction task: Observe the users' various approaches to solve the problem.)

Finally, subjects were given the opportunity to freely explore the system, describe problems with the MediaFinder interface, and offer suggestions for improvement.

A total of five subjects participated, one at a time. They were classified as sophisticated computer users in terms of computer usage experience, programming experience, and the number of applications they use on their machines. All of them were college students aged between 18 and 32 and their respective majors were Nutrition & Food Science, Hearing & Speech Science, Civil Engineering, Criminology, and Art.

All the subjects were able to complete all seven questions. The subjects quickly learned the concepts and were capable of using the semantic regions templates to answer the given tasks. They showed interest in the *fling-and-flock* metaphor and the concept of dynamic grouping of items.

In the verbal post survey, they were eager to use the *fling-and-flock* metaphor for other applications. One professional artist mentioned that he would like to use MediaFinder for organizing his paintings. He described taking photos of his paintings and that he wanted to group them based on categories such as color, painting medium, canvas size, timeline, and exhibition location. He said that it was crucial for him to show his buyers sample paintings based on their tastes and that he was eager to use MediaFinder as his painting organizing tool.

For the question concerning ease of use (scale 1–9) of the MediaFinder interface, they marked an average of 8.2 points and stated that the MediaFinder interface was not particularly difficult for them to use compared with other interfaces or the Windows file explorer. In previous validation studies well-liked systems got ratings in the range 6.0–7.5 while disliked systems got 3.0–5.0 [38]. Users in our study claimed that the first 10-min introduction was essential for them to understand the concepts and main features.

For the question about willingness to use (scale 1–9), they gave an average of 8.8 points under the condition that all the metadata are given and not annotated by themselves. This is a strong endorsement which is in harmony with user comments, but we are cautious in placing too much value on responses to such hypothetical questions for which there is no validation history.

Overall, the subjects were largely successful in grasping the cognitive concepts of semantic regions and using the MediaFinder interface, indicating satisfaction with the spatial organization and dynamic grouping strategies.

4.2. Evaluation of semantic regions construction

This second small study was intended to determine if users can learn to construct semantic regions, how difficult it is for users to construct them based on their conceptual models, and to identify cognitive trouble spots in the process.

Just as in the first usability study, a total of five subjects participated, one at a time. Three of them were computer science graduate students on campus and they were expert users of personal computers as well as MS Windows. They were highly experienced users in terms of computer usage experience, programming experience, and the number of applications they use on their machines. All the subjects already knew how to use semantic regions because they participated in the first user study either as pilot subjects or as actual subjects. This was intended to reduce the users' efforts to learn about semantic regions and how to construct them. Further studies with more junior users would be helpful to refine the interface and training methods.

The second usability study followed the same procedure as the first one. The background information survey showed that the experienced computer users frequently used their personal media data for entertainment such as playing movies, listening to music, and browsing photos. Just as in the first usability study, each subject was trained on MediaFinder for 10 min. After the completion of the practice tasks, they were given

4 tasks together with the test data set used for the first user study. The four tasks were as follows:

- Make a “Location” model identical to the given example
- 1. Make the same model as the sample location model you are given. Please confirm that your model works exactly the same as the given model by using the fling-and-flock metaphor.
- Make a “People” model without example
- 2. Choose any three people from the people list and construct regions for them. (Load their pictures as the background images of the created regions) And see how many pictures were taken for each person between 1995 and 2000. (Specify Conjunctive Semantics.)
- Combine models
- 3. Combine two models that you have created and see what states each person has been to, and find out the states that all three people have been to. (Combine the models horizontally.) How many pictures were taken for each person in California? How many pictures of Shneiderman were taken in each state? (Combine the models hierarchically.)
- Free construction
- 4. Make your own model with the given metadata such as date, people, location, name of conference, etc., and figure out any interesting facts from the created model.

Example models: seasonal calendar, yearly calendar, or monthly calendar.

All five subjects were able to complete the four tasks with occasional assistance from the administrator. Because the photos used for the user study had limited metadata (they only had location, people, date, conference name, and indoor/outdoor attributes) and the subjects were not familiar with the people in the photos, there was a clear limitation for the subjects to create some appealing models in a short time. However, these knowledgeable users could create their own models with MediaFinder after about 30 min of use. Subjects kept to simple Boolean combinations, with no subject creating a region that had more than two conjuncts.

In the verbal post-test survey, most subjects were enthusiastic about MediaFinder. There may have been social pressure to respond positively, since the subjects knew that the administrator of the experiment was also the developer of MediaFinder. On the other hand, studies of new interfaces often produce frustration for users. For the question about ease of use (scale 1–9) for construction tasks, they marked an average of 7.8 points, which was only a little lower than the first user study (8.2 points).

For the question about the willingness to use, they gave an average of 8.2 points, which was also a little bit lower than the first user study (8.8 points), under the condition that the metadata was given beforehand or automatically generated. Just as in the first user study, subjects were very concerned about the annotation process.

For the question about the helpfulness of the MediaFinder interface, they all gave high points (an average of 8 points). However, the average for the question about the preference of MediaFinder over other tools was lower (an average of 7 points) than those of previous questions because some of the subjects did not think that they needed a powerful tool for personal media management for the moment.

4.3. *Enhancement of the MediaFinder interface*

The two usability studies demonstrate that sophisticated users, with some training, are able to construct their own conceptual models for personal media management with

MediaFinder. The studies also helped to identify improvements that reduce the need for semantic specification and provide a semantics control panel.

The MediaFinder user interface has been redesigned to be similar to the Microsoft File Explorer to simplify learning how to specify region semantics and annotate personal media items.

The study on construction revealed that the two-step approach of creating new regions and then specifying semantics for each region was the primary difficulty for users because it took too much time and effort. Reducing the need to specify semantics in this manner would be a major benefit. While the capability for manually specifying semantics for a region empowers the expressiveness of semantics, shortcuts are possible for common simple situations. The MediaFinder user interface has been redesigned to support a simple drag-and-drop semantic regions construction to reduce the number of steps for region construction and semantics specification. Users can select a category tag from the semantic control panel and drag it onto the MediaFinder main panel, automatically creating a semantic region with the selected category tag. Then, all the personal media items bound to MediaFinder are dynamically regrouped according to the semantics specified in the newly created region. Users can spatially rearrange the newly created regions based on the specific tasks or their conceptual models.

5. Conclusion

Semantic regions is implemented in the MediaFinder user interface to allow users to rapidly explore and manage personal media items stored on their personal computers. Users can dynamically construct a variety of conceptual models and apply them to personal media items through the *fling-and-flock* metaphor. The *semantic regions* are spatially arranged on 2D space and used for personal media management tasks such as organization, meaning extraction, search, navigation, indexing, and distribution.

Two small usability studies on the MediaFinder interface revealed benefits, cognitive issues, and usability concerns. Sophisticated users succeeded in grasping the cognitive concept of semantic regions and using the MediaFinder interface. They expressed a preference for the spatial organization of information and dynamic regrouping of personal media items based on their conceptual models and specific tasks. This research on semantic regions makes four contributions:

- *Spatial organization*: *Semantic regions* enable users to specify spatial organization and dynamic reorganization of personal media based on their conceptual models. It extends the current system-oriented file management system to a user-oriented personal media management system by employing the semantics of personal media data. In addition, an interactive metaphor, *fling-and-flock*, that is designed and implemented for binding personal media items to a variety of users' conceptual models, has applicability for visual organization of numerous data domains including scientific and statistical data.
- *User interface*: The main design issues were ways of specifying the semantics for a region, showing the distribution of metadata collected from the personal media items, indexing personal media items, applying Boolean operations for search, and navigating the hierarchy of semantic regions. The visual design, dynamic aspects, and user controls implemented in *semantic regions* are also among its contributions to user-interface design.

- *Evaluation*: Although more work remains to be done in empirically characterizing the strengths of *semantic regions* as a personal media management mechanism, studies conducted so far have led to an increased understanding of its strengths. Further studies will attempt to refine this understanding, with the ultimate goal of generalizing results to apply to other 2D user-interface widgets.
- *Framework for extending the query model*: There are a number of possible extensions to *semantic regions* that might be used to increase the expressive power of the query language. Further work in this area will be needed to identify the potential extensions that are interesting and relevant to user tasks as well as realistically achievable.

Acknowledgments

We appreciate the comments on draft versions of this paper by Ben Bederson and other members of the Human–Computer Interaction Laboratory at the University of Maryland. Project sponsors included Ricoh, Intel, and Microsoft.

References

- [1] W. Arms, Glossary from 1999 Manuscript of Digital Libraries, MIT Press, Cambridge, MA, 2000.
- [2] N. Franklin, T. Federico, Representing described spatial and temporal situations in memory, in: Proceedings, AAAI-97, Workshop on Language and Space, Providence, RI, 1997, pp. 14–26.
- [3] G. Robertson, M. Czerwinski, K. Larson, D.C. Robbins, D. Thiel, M. van Dantzich, Data mountain: using spatial memory for document management, in: Proceedings of ACM UIST98, November 1998, pp. 153–162.
- [4] A. Cockbrun, B. Mckenzie, Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments, in: Proceedings of the ACM CHI2002 Conference, 2002, pp. 203–210.
- [5] R.N. Mayo, Reprint of the factoid webpage, Technical Note TN-60, Compaq Computer Corporation, Western Research Laboratory, July, 2001.
- [6] C.G. Bell, The cyber all project: a personal store for everything, Microsoft Research Technical Report MSR-2000-75, July 2000.
- [7] W.M. Newman, M.A. Eldridge, M.G. Lamming, Pepys: generating autobiographies by automatic tracking, in: Proceedings of ECSCW '91, Amsterdam, September 1991.
- [8] M. Lamming, Forget-me-not: intimate computing in support of human memory, in: FRIEND21, '94, International Symposium on Next Generation Human Interface, Meguro Gajoen, Japan, 1994.
- [9] S. Fertig, E. Freeman, D. Gelernter, Lifestreams: an alternative to the desktop metaphor, in: Proceedings of the ACM CHI96 Conference, New York, 1996, pp. 410–411.
- [10] J. Rekimoto, TimeScape: a time machine for the desktop environment, in: Proceedings of the ACM CHI '99 Conference, 1999.
- [11] B. Long, TimeStore: exploring time-based filing. Unpublished Paper, University of Toronto, 1994.
- [12] C. Plaisant, B. Milash, A. Rose, S. Widoff, B. Shneiderman, LifeLines: visualizing personal histories, in: Proceedings of the ACM CHI96 Conference, 1996, pp. 221–227.
- [13] H. Chen, S.T. Dumais, Bringing order to the web: automatically categorizing search results, in: Proceedings of the ACM CHI'00 Conference, 2000, pp. 145–152.
- [14] J.D. Novak, Concept maps and Vee diagrams: two metacognitive tools for science and mathematics education, *Instructional Science* 19 (1990) 29–52.
- [15] P. Zellweger, J. Mackinlay, L. Good, M. Stefik, P. Baudisch, City lights: contextual views in minimal space, in: Extended Abstract of CHI2003 ACM Conference on Human Factors in Computing Systems, 2003.
- [16] F. Shipman, R. Airhart, H. Hsieh, P. Maloor, J.M. Moore, D. Shah, Visual and spatial communication and task organization using the visual knowledge builder, in: Proceedings of the 2001 International ACM SIGGROUP Conference on Supporting GroupWork, 2001, pp. 260–269.
- [17] B.A. Nardi, S. Whittaker, E. Isaacs, M. Creech, J. Johnson, J. Hainsworth, Integrating communication and information through ContactMap, *Communications of the ACM* 45 (4) (2002) 89–95.

- [18] P. Dourish, W.K. Edwards, A. LaMarca, M. Salisbury, Presto: an experimental architecture for fluid interactive document spaces, *ACM Transactions on Computer–Human Interaction* 6 (2) (1999).
- [19] D.K. Gifford, P. Jouvelot, M.A. Sheldon, J.W. O’Toole, Jr., Semantic file systems, in: 13th ACM Symposium on Operating Systems Principles, October 1991.
- [20] E. Adar, D. Kargar, L.A. Stein, Haystack: per-user information environments. in: *Proceedings of the Eighth International Conference on Information and Knowledge Management*, ACM Press, New York, 1999, pp. 413–422.
- [21] D. Quan, K. Bakshi, D. Huynh, D.R. Karger, User interfaces for supporting multiple categorization, in: *Proceedings of Interact 2003*, 2003.
- [22] S. Card, J. Mackinlay, B. Shneiderman, *Readings in Information Visualization: Using Vision to Think*, Morgan Kaufmann, Los Altos, CA, 1999.
- [23] C. Ahlberg, B. Shneiderman, Visual Information seeking: tight coupling of dynamic query filters with starfield displays, in: *Proceedings of the ACM CHI ’94 Conference*, 1994, pp. 313–317.
- [24] C. Ahlberg, C. Williamson, B. Shneiderman, Dynamic queries for information exploration: an implementation and evaluation, in: *Proceedings of the ACM CHI ’92 Conference*, 1992, pp. 619–626.
- [25] C. Williamson, B. Shneiderman, The dynamic home finder: evaluating dynamic queries in a real-estate information exploration system, in: *Proceedings of ACM SIGIR ’92 Conference*, 1992, pp. 338–346.
- [26] B. Shneiderman, Inventing discovery tools: combining information visualization with data mining, in: *Proceedings of Discovery Science 2001*, Springer, Washington, DC, 2001.
- [27] S. Roth, P. Lucas, J. Senn, C. Gomberg, M. Burks, P. Stroffolino, J. Kolojejchick, C. Dunmire, Visage: a user interface environment for exploring information, in: *Proceedings of Information Visualization*, IEEE, 1996, pp. 3–12.
- [28] S.F. Roth, J. Kolojejchick, J. Mattis, M. Chuah, SageTools: an intelligent environment for sketching, browsing, and customizing data graphics, in: *Proceedings of the ACM CHI95 Conference*, 1995, pp. 409–410.
- [29] C. North, B. Shneiderman, Component-based, user-constructed, multiple-view visualization, in: *Extended Abstracts of CHI 2001*, 2001, pp. 201–202.
- [30] G.E. Krasner, S.T. Pope, A cookbook for using the model-view-controller user interface paradigm in smalltalk-80, *Journal of Object-Oriented Programming* 1 (3) (1988) 26–49.
- [31] H. Kang, B. Shneiderman, MediaFinder: an interface for dynamic personal media management with semantic regions, in: *Extended Abstracts of CHI 2003*, 2003, pp. 764–765.
- [32] S. Greene, E. Tanin, C. Plaisant, B. Shneiderman, R. Mushlin, L. Olsen, G. Major, S. Johns, The end of zero-hit queries: query previews for NASA’s global change master directory, *International Journal of Digital Libraries*, 2 ((2)(3)) (1999) 79–90.
- [33] M.C. Chuah, Dynamic aggregation with circular visual designs, in: *Proceedings of the IEEE Symposium on Information Visualization*, October 1998, pp. 35–43.
- [34] J. Goldstein, S.F. Roth, Using aggregation and dynamic queries for exploring large data sets, in: *Proceedings of the ACM CHI ’94 Conference*, 1994, pp. 23–29.
- [35] L.Tang, Dynamic aggregation, CMSC838 class project in spring 2001, http://www.cs.umd.edu/class/spring2001/cmssc838b/Project/Lida_Tang/index.html
- [36] B. Shneiderman, H. Kang, B. Kules, C. Plaisant, A. Rose, R. Rucheir, A photo history of SIGCHI: evolution of design from personal to public, *ACM Interactions* 9 (3) (2002) 17–23.
- [37] B. Shneiderman, H. Kang, Direct annotation: a drag-and-drop strategy for labeling photos, in: *Proceedings of International Conference Information Visualisation (IV2000)*, 2000.
- [38] J.P. Chin, V.A. Diehl, K.L. Norman, Development of an instrument measuring user satisfaction of the human–computer interface, in: *Proceedings of ACM CHI’88 Conference*, New York, 1988, pp. 213–218.
- [39] G. Bell, A personal digital store, *Communications of the ACM* 44 (1) (2001) 86–91.
- [40] G. Bell, J. Gray, Digital immortality, *Communications of the ACM* 44 (3) (2001) 28–30.
- [41] R. Mander, G. Salomon, Y.Y. Wong, A pile metaphor for supporting casual organization of information, in: *Proceedings of the ACM CHI92 Conference*, 1992, pp. 627–634.
- [42] G.E. Moore, *Gigabits and Gigabucks*, University Video Corporation Distinguished Lecture, 1996.