TECHNICAL RESEARCH REPORT

Exploring Personal Media: A Spatial Interface Supporting User-Defined Semantic Regions (2004)

by Hyunmo Kang, Ben Shneiderman

TR 2005-51

ISR
INSTITUTE FOR SYSTEMS RESEARCH

ISR develops, applies and teaches advanced methodologies of design and analysis to solve complex, hierarchical, heterogeneous and dynamic problems of engineering technology and systems for industry and government.

ISR is a permanent institute of the University of Maryland, within the Glenn L. Martin Institute of Technology/A. James Clark School of Engineering. It is a National Science Foundation Engineering Research Center.

Web site http://www.isr.umd.edu

EXPLORING PERSONAL MEDIA: A SPATIAL INTERFACE SUPPORTING USER-DEFINED SEMANTIC REGIONS

Hyunmo Kang*, Ben Shneiderman

{kang, ben}@cs.umd.edu

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

Abstract

Graphical mechanisms for spatially organizing personal media data could enable users to fruitfully apply their mental models. This paper introduces *Semantic Regions*, an innovative way for users to construct mental models by drawing regions on 2D space and specifying the semantics for each region. Then users can apply personal ontologies to personal media data using the *fling-and-flock* metaphor. This allows photos (or other personal media) to be dragged to the display and automatically grouped according to time, geography, family trees, groups of friends, or other conceptual maps. The prototype implementation for *Semantic Regions*, MediaFinder, was validated with two usability tests for usage and construction of user-defined mental models. Applications and extensions are discussed.

1. Introduction

Computer users deal with large numbers of personal media objects such as images, audio clips, voice mail, video clips, web pages, emails, and document files.

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

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

There is no formal definition of "personal media" but we use it to describe computerized information of familiar items, such as family photos, and strong attention to browsing. This vast array of personal media objects presents challenges for most users, who often struggle to interpret, explore, arrange, and use personal media objects. They wrestle with 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 models of it in their minds. There is no unique or right model. Rather, the mental model is personal, has meaning for the individual who creates it, and is tied to a specific task. Even in a simple personal photo library, images can be organized by timelines, locations, events, people, etc. depending on users' mental models. Despite the diversity of users' mental 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 and implementing a novel user interface for exploring personal media is to make it easier for end-users to construct, represent, and apply their mental models to personal media data for more flexible and task-specific personal media management.

Another challenge in designing a personal media explorer 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 can be a solution because they allow rapid information retrieval in 2D graphical user interfaces. Several experiments (Robertson et al., 1998; Cockbrun and Mckenzie, 2002) 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 the semantics of well-known display representations or users' preferred mental models, it is expected to improve task performance even more.

Numerous display representations (e.g. diagrams and charts) are used for showing information because they often make data more comprehensible and ensure faster data perception. Familiar representations used for displaying information include calendars, timelines, maps, organization charts, floor plans, body diagrams, building layouts, integrated circuit block diagrams, periodic tables, and so on (**Figure 1**). Unlike the standard mathematical or statistical visualizations such as histograms, line graphs, pie charts, scatter plots, lists, tables, grids, etc., 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 relationship. 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

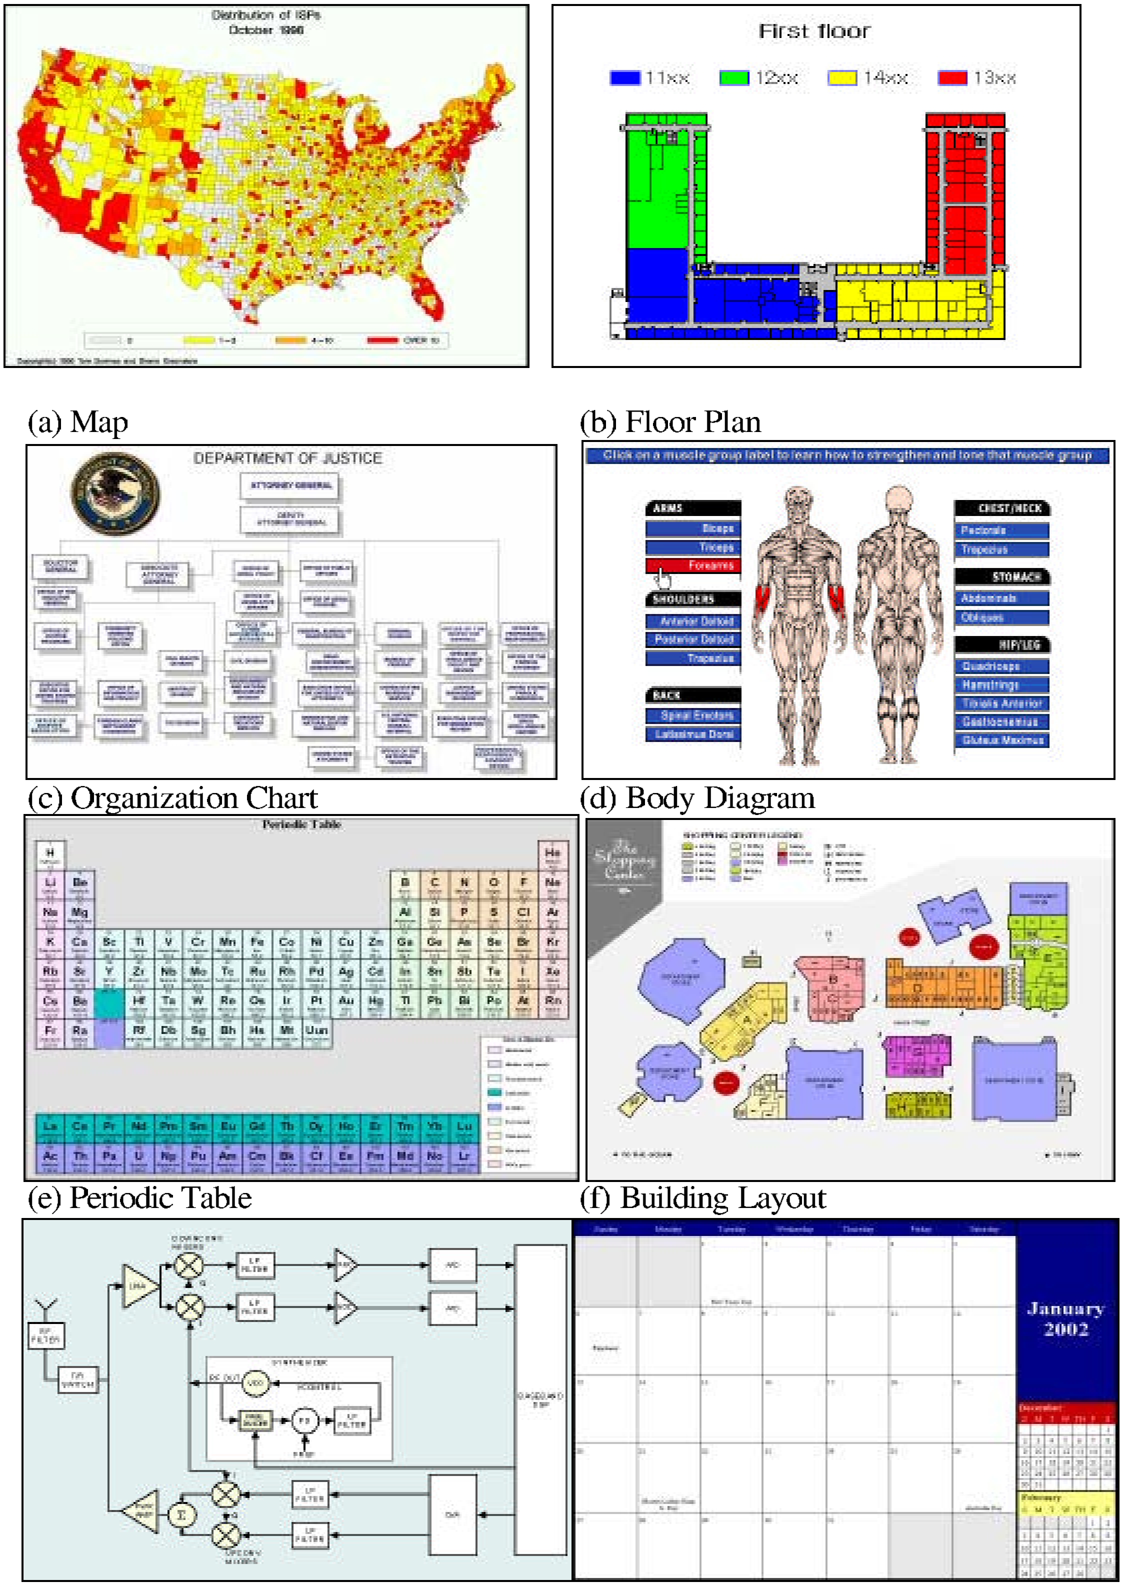

(a) Map (b) Floor Plan

(c) Organization Chart (d) Body Diagram

(e) Periodic Table (f) Building Layout

(g) Integrated Circuit Block Diagram (g) Calendar

Figure 1 Display representation examples 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 in the museum, 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 well as personal information can be combined with visually compelling display representation to provide engaging perspectives for different information tasks.

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, ranging from processor speed, storage capacity, data transmission rate, scanner speed, camera resolution, audio/video encoding, printing/display quality, etc., there have been a great deal of research efforts made concerning Personal Information Storage.

The Factoid project (Mayo, 2001) at Compaq Research Lab imagined small ubiquitous devices that would capture many small facts per day and which are each only a few hundred bytes in size. Based on 1000 facts a day, they calculated 73MB/year which can thus be saved forever and constitute a sort of history of the users' life. In the CyberAll (Bell, 2000) project at Microsoft, Famed researcher, Gordon Bell 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.

Other projects such as those at XRCE (Xerox Research Center Europe) Cambridge focused on augmenting memory. Pepys 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" (Lamming , 1994), information is automatically collected about a user's activities such as telephone calls made, e-mail notes received, and meetings attended. This information is then stored and presented as a personal biography with activities presented in time order. The biography can be browsed, and it may be filtered to show only a particular type of activity.

2.2 Information Organization and Retrieval

Due to the progress in technology and a corresponding desire by users to digitize and store their personal information, the amount of personal media data on individual machines is increasing rapidly. 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.

The LifeStreams (Fertig et al., 1996) system 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, a standard single-inheritance file hierarchy, difficulty in making use of archived information, difficulty in getting a big picture view, and difficulty in managing schedules and reminders. TimeScape (Rekimoto, 1999) shows an interesting desktop metaphor called TMC (Time Machine Computing). The 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 (Long, 1994) and LifeLines (Plaisant et al., 1996) are other time-based information indexing and retrieval systems. The designers of Stuff I've Seen (Chen and Dumais, 2000), 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 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. A common strategy was to use a timeline as a new organizing metaphor for the document management system, to overcome some of the problems of the rigid folder hierarchy. 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 metaphor is still fixed and rigid and cannot be applied to the problems that require multiple perspectives of the data.

On the other hand, there have been several research attempts at using different organizing metaphors than a timeline. Richard et al. (1992) designed an interesting desktop user interface metaphor, Pile, to support the casual organization of documents. The interface design based on the pile concept suggested uses of content awareness for describing, organizing, and filing textual documents. Data Mountain (Robertson et al., 1998) allows users to place documents at arbitrary positions on an inclined plane in a 2.5D desktop virtual environment using a simple 2D interaction technique. This user interface was designed specifically to take advantage of human spatial memory in retrieving a document from a virtual desktop working space. Tinderbox (http://www.eastgate.com/Tinderbox), Niagara (Zellweger et al., 2001), Visual Knowledge Builder (Shipman et al., 2001) are visual information workspaces supporting 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 (Nardi et al., 2002) 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. This research is based on the fact that people invest considerable effort in maintaining links with networks of colleagues, acquaintances and friends, and that these networks are a significant organizing principle for work and information. Those tools took advantage of human spatial memory in retrieving documents from a virtual desktop working space. Although spatial organization of information enables users to retrieve documents quickly, it is not easy for users to represent multiple organizations based on their tasks or mental models. Furthermore, there is no good way to represent an organizational hierarchy with these tools.

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 (Dourish et al., 1999) 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 (Gifford et al., 1991) is an information storage system that provides flexible associative access to the system's contents by automatically extracting attributes from files with file type specific transducers. In this system, 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 (Adar et al., 1999; Quan et al., 2003) 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 various views in order to handle the different types of information, different visualizations, and multiple perspectives for different aspects of an information task (Card et al., 1999; Nardi et al., 2002). However, it is difficult for users to find generalized visualization tools satisfying their demands, or to change visualizations dynamically according to their mental models or specific tasks.

Spotfire (http://www.spotfire.com) is one of the most successful generalized information visualization and visual data-mining tools. Easy import/export of data, rapid change to axes, color coding or size coding, and direct data manipulation (Ahlberg et al., 1992; Ahlberg and Shneiderman, 1994, Williamson and Shneiderman, 1992) have made it a leading generalized data visualization tool. 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 map, calendar, genealogy trees, etc. Like other successful statistical data mining tools such as SAS JMP, DataDesk, and SPSS Diamond (Shneiderman, 2001), Spotfire is more appropriate for visualizing scientific and quantitative data rather than metadata-based personal media data.

Visage (Roth et al., 1996) 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 (Roth et al., 1995) 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 (North and Shneiderman, 2001) 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.

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) (Krasner and Pope, 1988) paradigm, the user input, the modeling of the external world, and the visual feedback to the user 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' mental models and personal ontologies.

3. A Spatial User Interface Supporting User-defined Mental Models

To advance research on personal media, we propose *Semantic Regions,* based on one major hypothesis: Spatially organized information based on the semantics of personal media and the users' mental models for managing personal media will greatly improve task performance as well as user satisfaction. Our MediaFinder prototype has been implemented 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 mental models in their minds. 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 Christian, but others may break it down into more detailed categories such as rock, jazz, R&B, rap, country, classical, Latin, Christian, and so on. Even in the classical category, some may conceive of music by historical periods such as Renaissance, Baroque, Classical, Romantic, 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 mental maps can be differently organized and represented on a 2D space. For example, some people may see R&B as closer to rock, whereas others see 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 R&B, and four times as large as jazz). They may even make use of shapes or connecting lines to make their mental maps to be more understandable and recognizable.

For another example, people may have a rough map of their neighbors. Unlike an actual map, it can represent a personal cognitive map of strength of interest in a local park, a neighbor's house, a grocery store, a public swimming pool, library, parking lot, 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 their information on their own way.

There is no right or unique mental model for personal media data and because of this diversity, 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 mental models and specific tasks.

3.2 Semantic Regions: Interactive Personal Media Query Widget

Semantic Regions are query regions drawn directly on a two-dimensional information space. Users can specify the shapes, sizes, and positions of the regions in two-dimensional space and thus form the layout of the regions meaningful to them *Semantic Regions* created with MediaFinder (see examples in **Figure 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 (**Figure 2**(a)). Once the regions are created and arranged, users can specify semantics for each 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.

(a) World map model

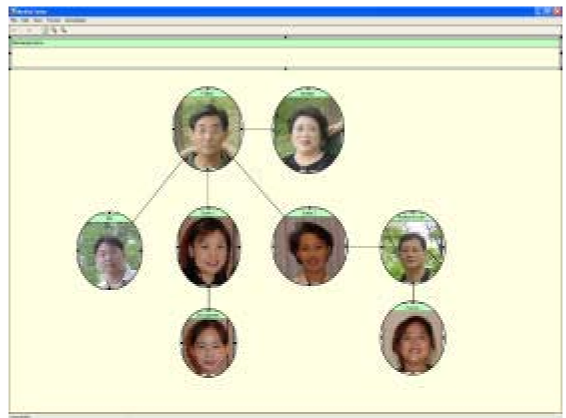

(b) Genealogy tree model

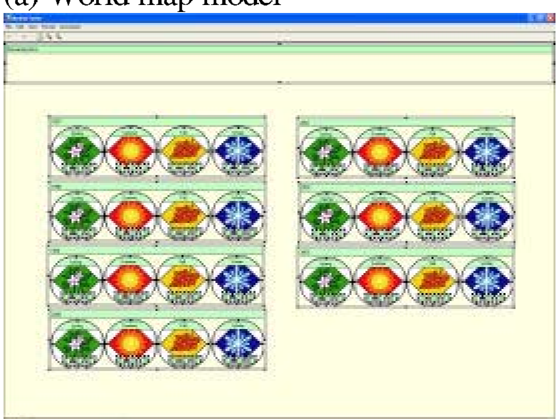

(c) Season calendar model

(d) Group organization chart model

Figure 2 Example models of *Semantic Regions* constructed with MediaFinder

In **Figure 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 regions, 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.

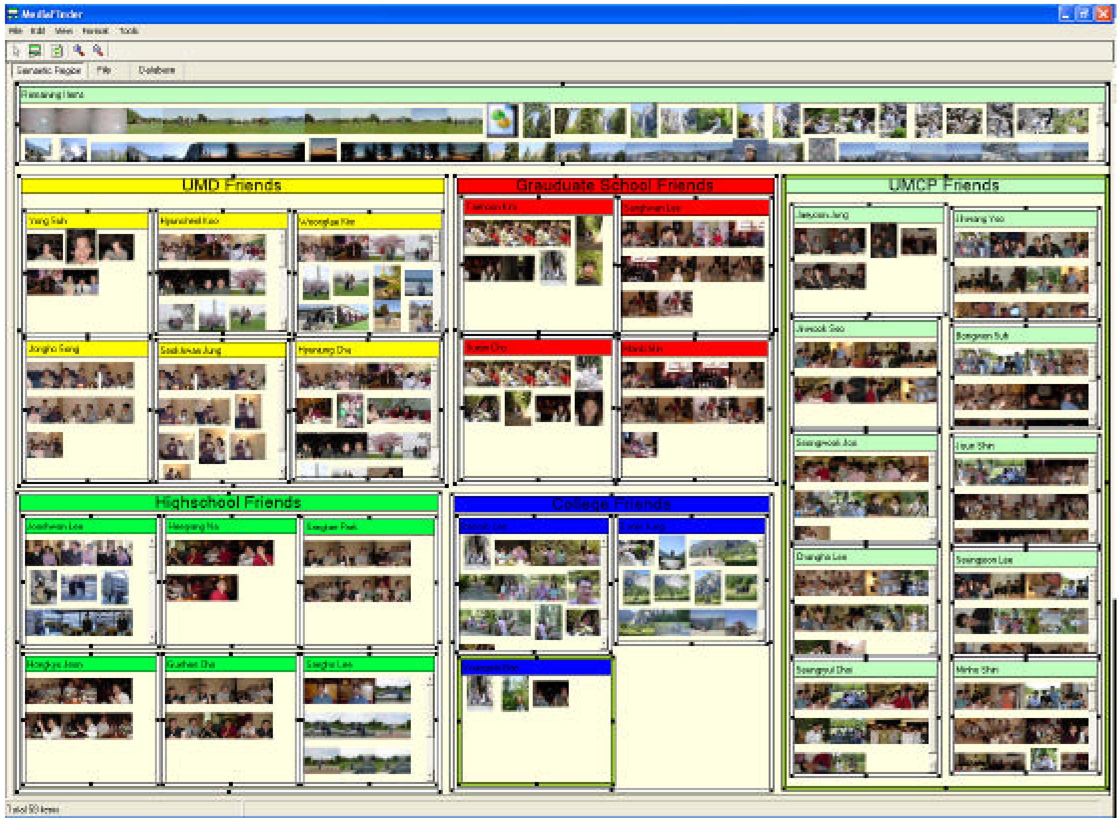

Figure 3 A friend group mental model: Each region represents a person and contains all the photos annotated with the name defined in it. The regions are grouped into 5 clusters to represent different friend groups (UMD friends, high school friends, graduate school friends, college friends, and UMCP friends). Each group has its own color to represent the different group of friends.

A semantic region works just like a folder. 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 its 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 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. Unlike the folder structure, semantic regions need not to be mutually exclusive, and thus enable more flexible organizations of personal media data.

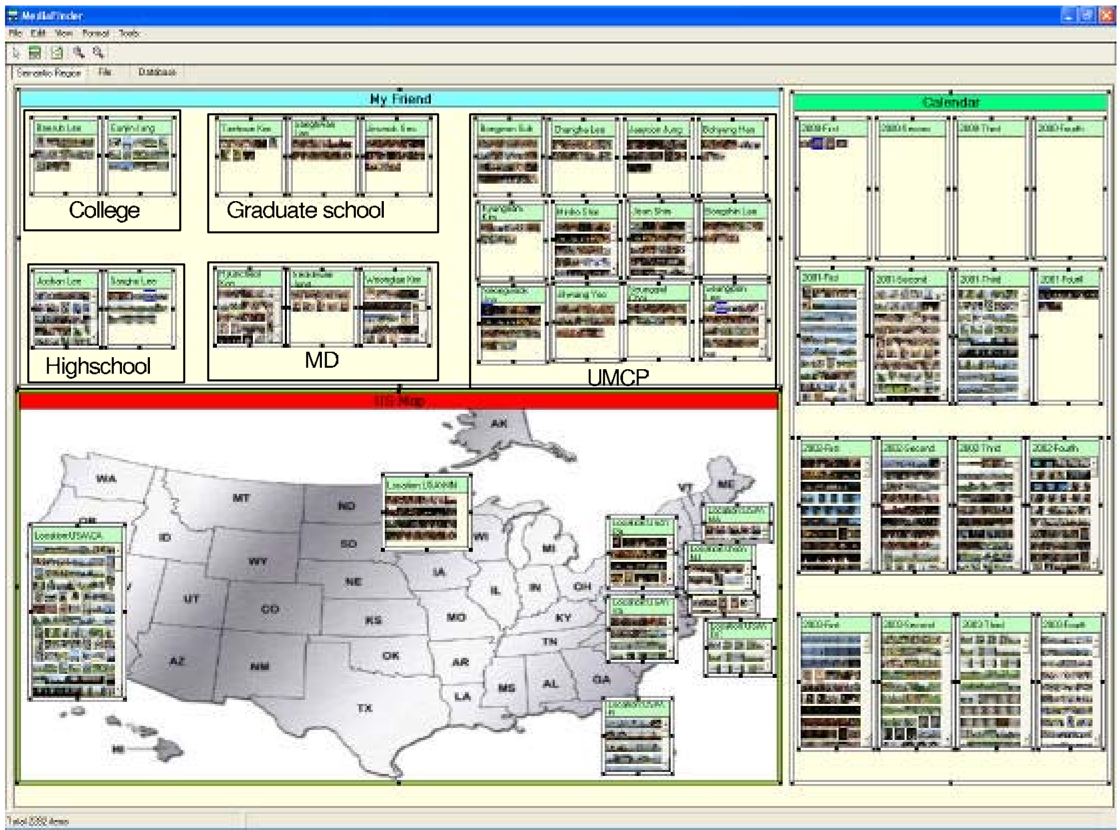

Figure 4 Three mental models, friend groups, quarter year calendar, and US map (clockwise from the top region) are combined together. The regions in the My Friend region are grouped into 5 clusters to represent different friends groups (high school friends, college friends, graduate school friends, UMD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

Using *Semantic Regions*, users can combine a few heterogeneous mental models simultaneously to observe the relationships among them. In **Figure 4**, three mental models, the US map, the year calendar, and my friends group are combined together and bound to photos through the *fling-and-flock* metaphor.

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 mental 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 two-dimensional workspace and many functions that help users construct and use *Semantic Regions*. In the workspace, users can freely create, remove, resize, move, and change the z-order of *Semantic Regions* just like object-oriented drawing tools. Several drawing and formatting features implemented in MediaFinder enable users to easily outline their mental models on 2D media space. After users sketch their mental models, they can specify the semantics for each region to represent their meanings.

3.4.1 Semantic Regions Creation

The creation of *Semantic Regions* is straightforward: users simply 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. Users can also specify a title for a region for the same purpose. The child regions can be constructed recursively inside a region. The background of each region is used as a new 2D workspace on which new child regions are created. 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 various perspectives of the contained personal media item. Each region supports five basic views (list, small icon, large icon, details, and thumbnails) just as the Windows file explorer does. Users can launch the associated applications by double-clicking over a personal media item in these views just as they can do with Windows file explorer. Besides the five basic views, three additional views (aggregation, intersection, and background) were implemented to provide a quick overview of the model as well as to represent the interrelationships of the 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 (**Figure 5**(a)). The background view enables users to easily remember the implication of the region through a background image. It also helps users perceive the internal hierarchy of a region through the overview of the child regions (**Figure 5**(b)). Additionally the intersection view was designed to represent the amount of personal media items shared among the regions.

The overall design of multiple views for *Semantic Regions* was based on the well-known data visualization strategy, "Overview first, zoom-in, then filter, and detail-on-demand" (Card et al., 1999). Users can overview the model and distribution of personal media data first, and then see more details by maximizing the region. The design was also based on the theory of query preview (Greene et al., 1999) or dynamic aggregation (Chuah, 1998; Goldstein and Roth, 1994; Tang, 2001).

Users can draw lines among 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' mental models.

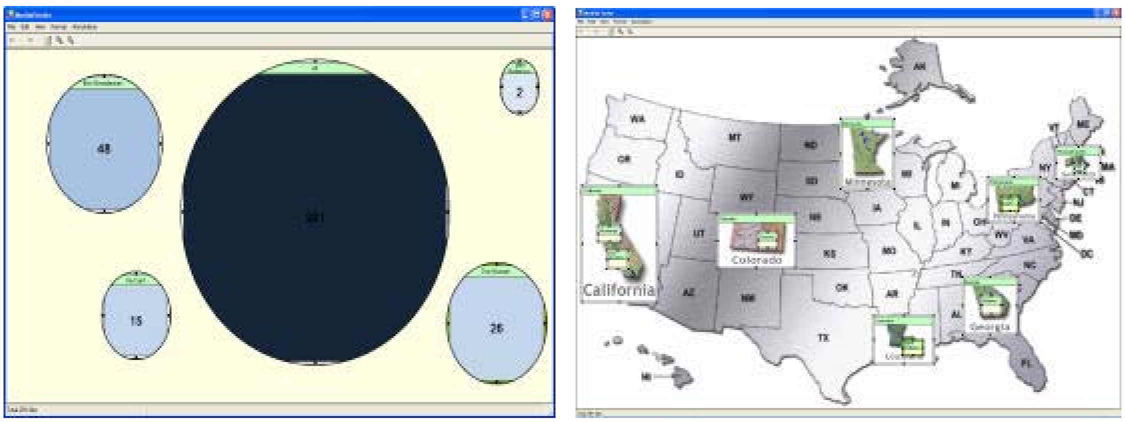

(a) Aggregation view with size and color encoding

(b) Background view with background image and overview of child regions

Figure 5 Advanced views of *Semantic Regions*

3.4.2 Specification of Semantics

Each semantic region should have semantics to enclose the personal media items. The semantics of a region consist of 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 depending on the types of personal media data. These different types of attributes are joined conjunctively to form the semantics of a region. Right-click on a region pops up a menu. From this menu, users can select the "semantics menu", which brings up the *Semantic Regions Attribute dialog box*. The Semantic Regions Attributes dialog box is composed of five tabs: name, file, date, people, and category. The name tab is used for specifying the title of the region. In the file and date tabs, users can make use of a variety of file attributes of the Windows file system to specify the semantics for a region. Once the attribute values are specified, all the personal media items that do not satisfy the criteria will be filtered out dynamically from the region.

On the other hand, users can define their own category attributes and use them in specifying the semantics for a region. They can even define multiple categories and create a hierarchy for each category. For example, users can define both a location category that has a three-level (country-state-city) hierarchy and a season category that has just one-level hierarchy to specify the user-defined semantics for a region.

In the Semantic Regions Attributes dialog box, either the category tab or the people tab is used for specifying the user-defined attributes for a region. In the category tab, a user-defined attribute is specified by selecting a node from the category tree and adding it to the selected category listbox (right window) (**Figure 6**(a)). The selected node is converted to a string like a full path name of a folder in the windows file system. It starts with the category name that is followed by a colon and the series of ancestor node names separated by a back-slash character (e.g. location:USA\MD\College park). The specified categories can be removed from the selected category listbox by pressing the left arrow button. Users can create, modify, or remove a category node by using the "new", "edit", and "delete" buttons located at the bottom of the category hierarchy window. Although the "People" category is one of many user-defined categories, MediaFinder was designed to have a separate "People" tab in the Semantic Regions Attributes dialog box (**Figure 6**(a)) because the names of people are frequently used as key metadata in personal media data. As in the category tab, users can create, modify, and remove a name from the People listbox through “new”, “edit”, and “delete” buttons in the people tab.

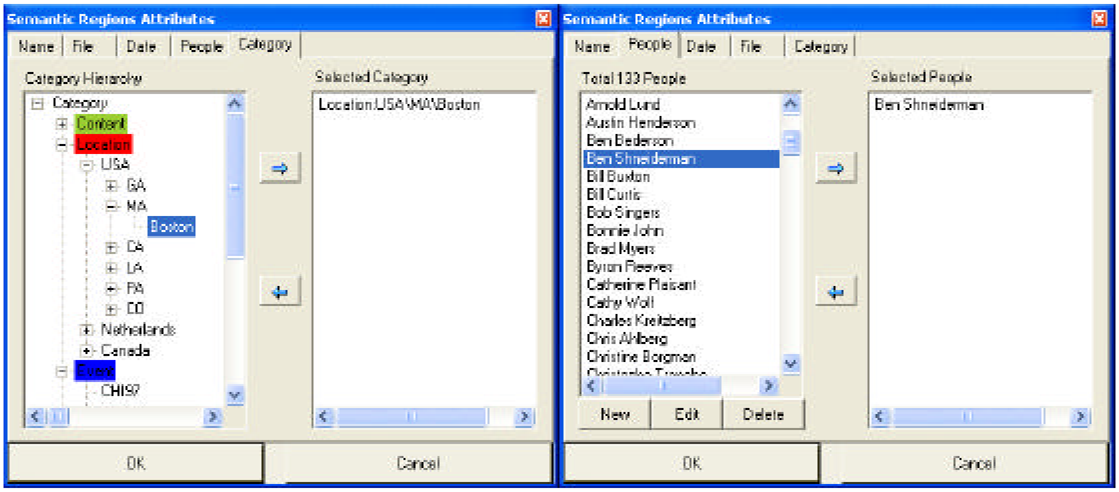

(a) Specification of category attributes (b) Specification of people attribute

Figure 6 Specifying the user-defined attributes for a region in the category and the people tab of the Semantic Regions Attribute dialog box.

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

Personal media items can be annotated with user-defined attribute values such as category and people. As mentioned above, a user-defined attribute value is a string that starts with the name of a category and a colon, and is followed by a series of the names of sub-categories separated by a back-slash character. 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 such as category and people 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 to represent in the Semantic Regions Attributes dialog box. These pre-constructed hierarchies of the user-defined attributes are used for specifying the semantics for a region.

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 various different mental models to the same personal media data set by opening the previously saved files. Also, users can import and combine multiple heterogeneous models together in MediaFinder to observe the interrelationships among the models.

The models can be combined horizontally or vertically (**Figure 7**). 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 interrelationships 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.

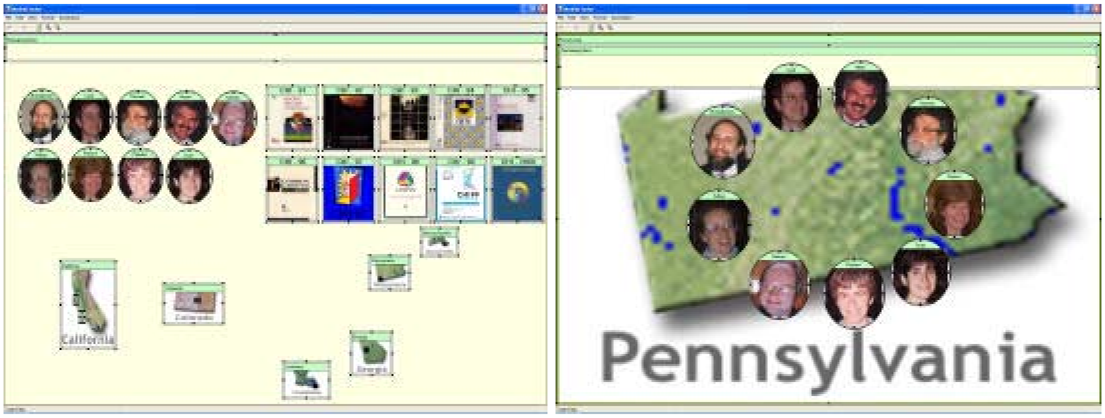

(a) Three models combined horizontally (b) People model inside a location region

Figure 7 Multiple models can be combined horizontally or vertically

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

3.5 Using *Semantic Regions* in MediaFinder

Binding a semantic regions model with the personal media items enables users to use MediaFinder as a tool for managing personal media data such as organization, meaning extraction, search, navigation, indexing, and distribution. The following sections describe how MediaFinder was designed and implemented to deal with each personal media management task.

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 (**Figure 8**). The remaining items region is a specially designed semantic region. Unlike other semantic regions, it is not allowed to have any semantics so that it can contain all the personal media items that do not satisfy any of the criteria for other regions. Likewise, each region was designed to have a child remaining items region to contain all the personal media items that do not satisfy any of the criteria of the child regions but which are contained in the parent region.

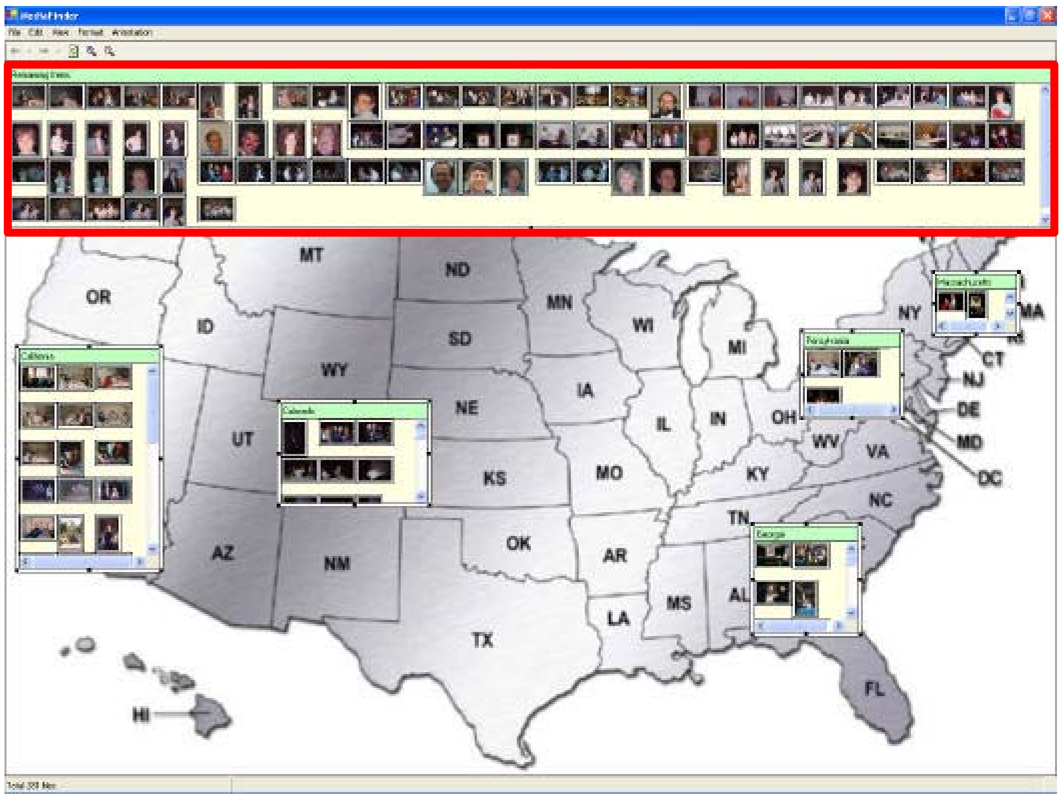

Figure 8 The photos that do not satisfy the semantics of five state regions are collected in the *Remaining Items* region (red rectangle) located at the top of MediaFinder main window.

During *fling-and-flock*, MediaFinder collects all the attribute values of the bound personal media items and keeps them in the Semantic Regions Attributes dialog box. 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 like query preview (Greene et al., 1999). The collected user-defined attributes

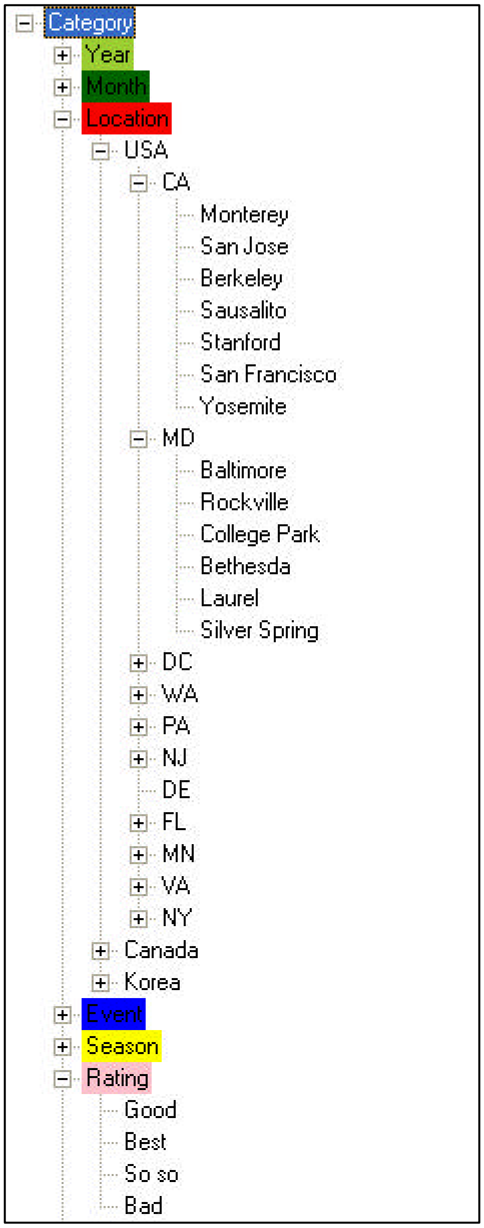

Figure 8 User-defined attributes collected by MediaFinder through fling-and-flock

are translated to form a semantic hierarchy as shown on the right. **Figure 8** 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 (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. In other words, the intersection of any two regions in MediaFinder does not have to be empty. This design is based on the fact that users' mental 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 mental models and observe the interrelationship among them. The multiple inclusion of a personal media item across the regions often depicts both the intra-relationships and the inter-relationships of the different models.

3.5.3 Navigation

Semantic regions can be maximized to show the detail information of the contained personal media items or to represent the hierarchy of the child regions (**Figure 9**). **Figure 9**(a) shows the USA map model that contains five state regions. A region is maximized to fill the parent region with animation when it is double-clicked. **Figure 9**(b) shows the maximized 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. In **Figure 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 maximized recursively and shows the contained items with a thumbnail view (**Figure 9**(d)). The titles of the entire maximized regions are displayed contiguously to represent the current level in the hierarchy (**Figure 9**(d)). The maximized region shrinks back to its original size when it is double-clicked again. All of the maximizing and minimizing processes are animated to reduce the users' confusion that might be caused by the change of region hierarchy.

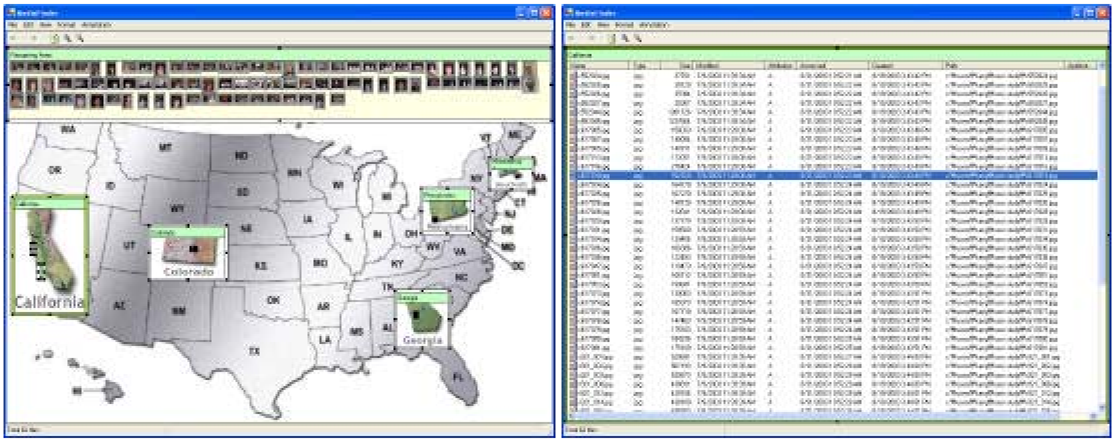

(a) Maximize California by double-click (b) Detail content list in California region

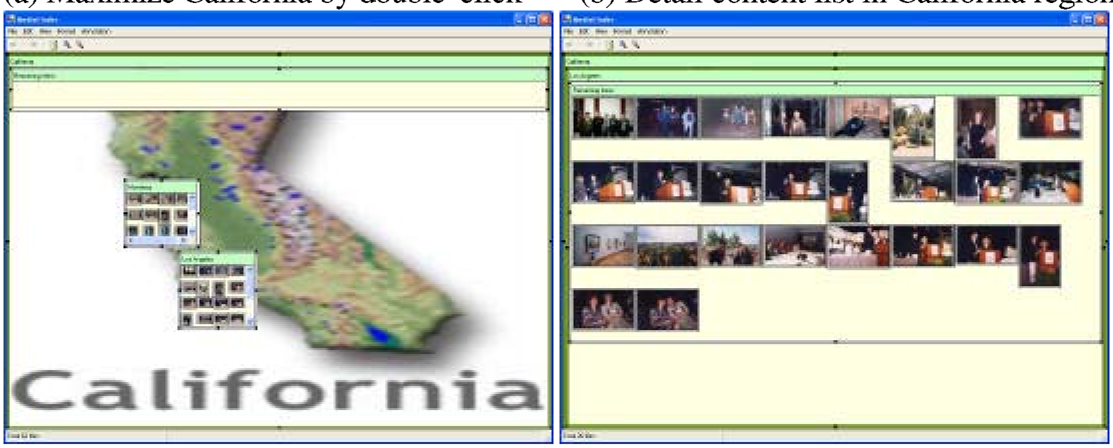

(c) Two city child regions in California (d) Maximize a city region recursively

Figure 9 The maximized region shows more details about the contained personal media items or represents the hierarchy of child regions.

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. **Figure 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 to be easily identified (**Figure 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 (**Figure 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 (**Figure 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 **Figure 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.

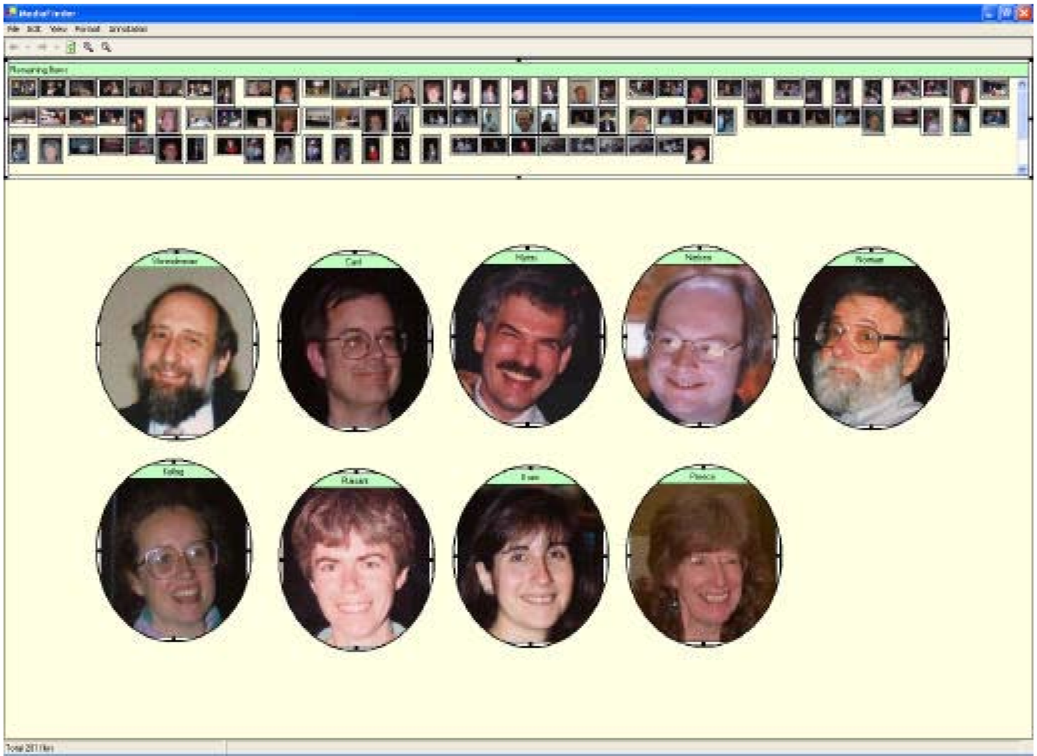

(a) “HCI researchers” model bound to the conference photos

(b) The selected item and all the containing regions are highlighted red

Figure 11 Region brushing: an interaction technique to visualize the shared items among semantic regions (Continued)

(c) Regions containing any of the items in the mouse-overed region are highlighted

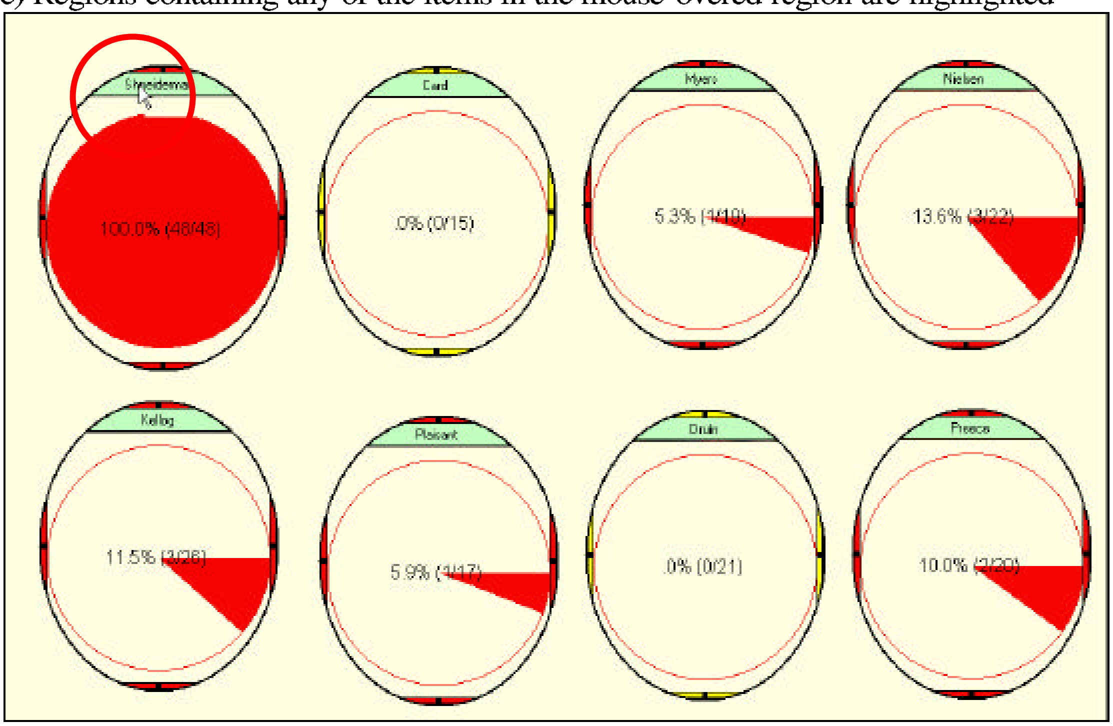

(d) Intersection view to represent the number of shared items among the regions

Figure 10 Region brushing: an interaction technique to visualize the shared items among semantic regions.

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 that 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.

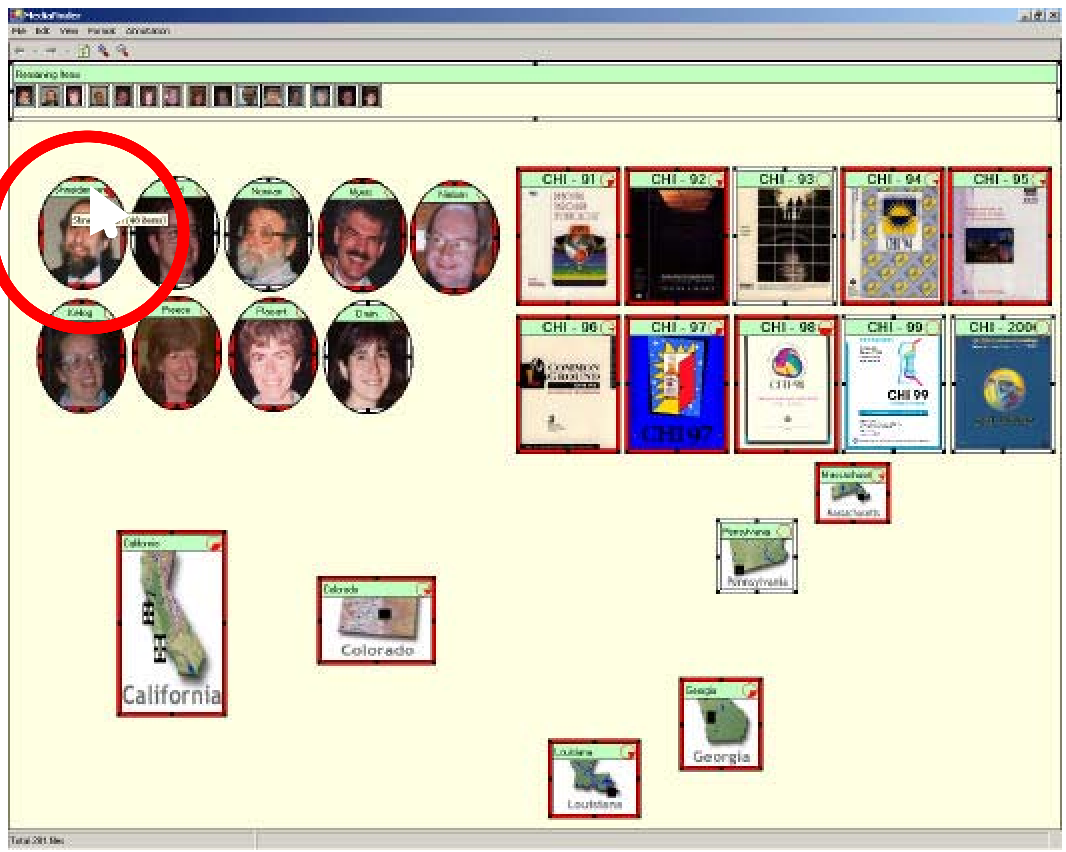

Figure 11 Extracting meanings concerning the interrelationships of the multiple models through the region brushing technique. By placing the mouse on Shneiderman's region, users can see which conference and which state his photos were taken.

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.* Boolean AND or OR operations are applied to multiple regions but the NOT operation is not allowed to be applied to more than one region since NOT is a unary operator. 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. 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 the Boolean regions (the regions created by applying Boolean operations). With the Boolean operations, users can specify any logical combination of the semantics for a region.

Figure 12 shows the process of applying Boolean operations to the regions. Suppose an example, "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 **Figure 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 (**Figure 12**(b)). The third step is to select both the Ben Shneiderman region and the Jenny Preece region to apply an OR operation (**Figure 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 (**Figure 12**(d)). Finally, the AND operation is applied to the NOT and the OR regions (**Figure 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))).

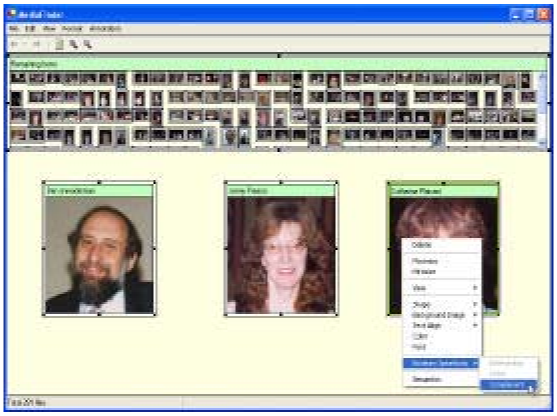

(a) Step 1: Apply NOT to the region of Catherine Plaisant

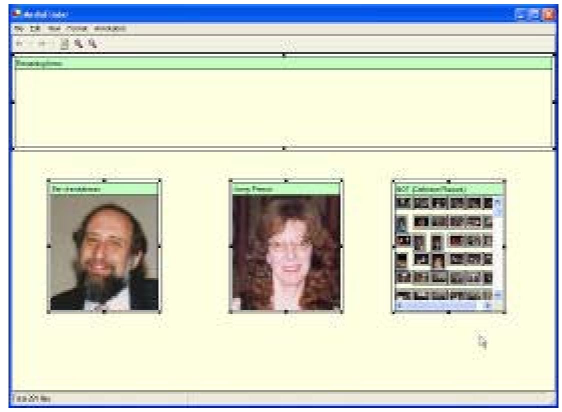

(b) Step 2: A new NOT region is created over the region

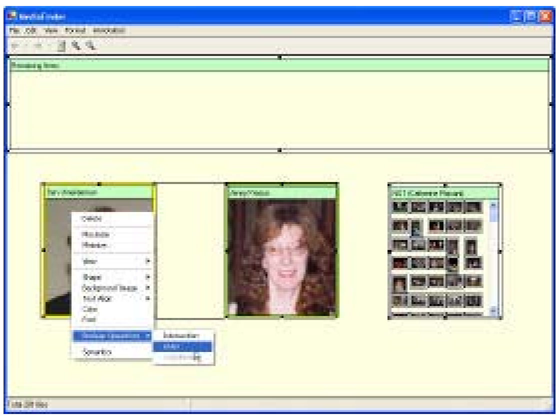

(c) Step 3: Apply OR to Ben Shneiderman and Jenny Preece regions.

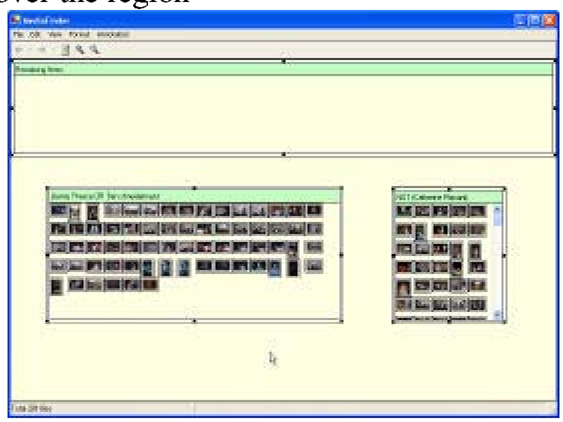

(d) Step 4: A new OR region is created over the regions

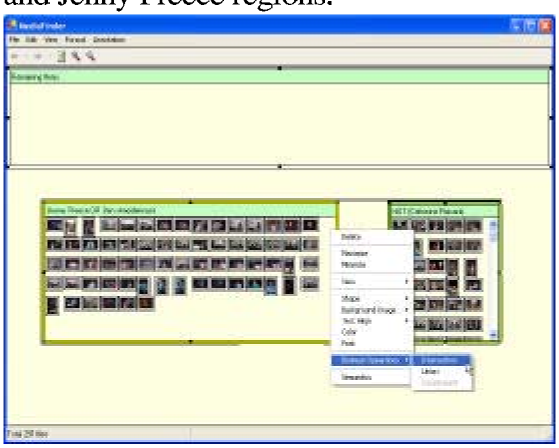

(e) Step 5: Apply AND to two Boolean regions

(f) Step 6: Final AND region constructed over the OR and the NOT regions

Figure 12 The steps for applying Boolean operations to get the Boolean region of (NOT (Catherine Plaisant)) AND ((Ben shneiderman) OR (Jenny Preece)))

Boolean 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 Boolean regions that use this region as an operand would be updated together. Recursively, the content change of the Boolean regions may also affect other Boolean regions that use the updated Boolean regions as their operand regions. Therefore, the Boolean regions are updated successively based on the relationships with their operand regions and thus the order of Boolean operations forms a hierarchy. Cycles are not allowed in the hierarchy of Boolean operations.

3.5.6 Indexing and Distribution

Semantic Regions can be used for indexing personal media items. If users select personal media items in one region and drag them to another region, the attribute values of the selected items are updated with those specified in the dragged region. If the semantics of the dragged region and the semantics of 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 dragged region conflict with the semantics of the original region, the personal media items are removed from the original region and added to the dragged region.

In MediaFinder, the drag-and-drop metaphor is used for indexing personal media items just as it is used in direct annotation (Shneiderman and Kang, 2000). In **Figure 13** (a), five photo items in the remaining items region are selected and dragged onto the Minnesota region. If the items are dropped on the region, they are annotated with the location attribute value "location:USA\Minnesota" that is specified in the

Minnesota region. Subsequently, the dragged items are removed from the remaining items region and added to the Minnesota region because they now satisfy the semantics of the region (**Figure 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.

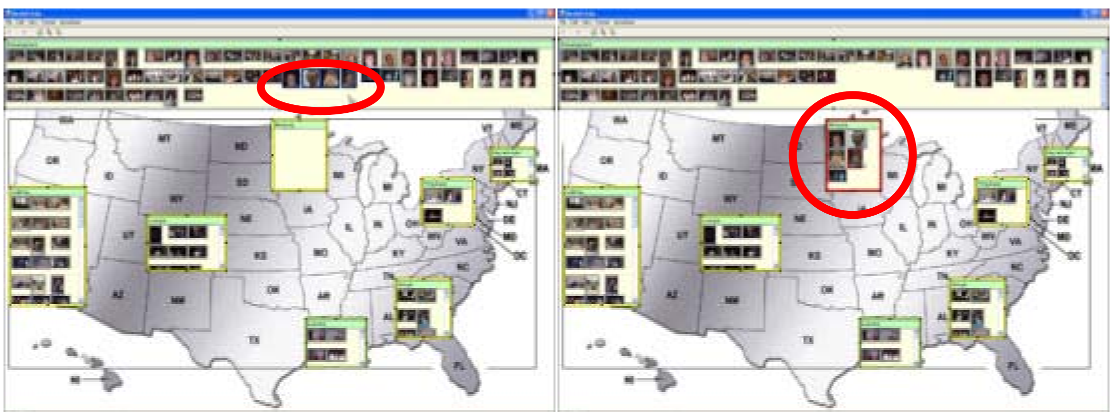

(a) 5 photos are dragged from remaining items region onto the Minnesota region

(b) Dragged items are annotated with the semantics of the Minnesota region

Figure 13 Indexing the personal media items by dragging them onto the semantic region

4. Evaluation

Studying the use of semantic regions is important for three reasons: To evaluate the usability and benefits of MediaFinder (in terms of performance and user satisfaction), to discover potential user interface improvements, and to gain a deeper level of understanding about users' ability to understand, construct, and operate semantic regions strategies. Two 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 mental models.

4.1 Evaluation of Semantic Regions Use

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

Then, when confident to continue, each subject was given a set of practice tasks, and allowed to ask any questions about MediaFinder interface. The test data set consisted of 281 photos that were taken at the 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 PhotoFinder4 (http://www.cs.umd.edu/photolib). In addition to photos, four semantic regions templates were given: Location Map (7 States and 8 Cities in US), Conference Calendar (10 Years of Calendar), People (9 Most frequently appearing people), Combined Model (Combination of the above three models). A total of 7 revised 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 doesn't 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 mental 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 mental models)
6. What was the name of the conference held in Atlanta, Georgia? Find out who didn't 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 novice 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 18~32 and their respective majors were Nutrition & Food Science, Hearing & Speech Science, Civil Engineering, Criminology, and Art.

All the subjects were able to complete all 7 questions. The subjects were quick to learn the concepts and usage 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 media items.

In the verbal post survey, all the subjects showed their interests and 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 dynamically. He said that he had a lot of paintings and always took pictures of them for future display and that he wanted to group them based on various categories such as color, painting medium, canvas size, timeline, exhibition location, and so on. He said that it was crucial for him to show his buyers sample paintings conveniently based on their tastes and that he was eager to use MediaFinder as his painting organizing tool.

For the questionnaire concerning ease of use (scale 1 to 10) 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. They all agreed that the first 10-minute introduction helped them understand the concept and main features, but without it, they would not have been able to answer the questions. For the questionnaire about the willingness to use (scale 1 to 10), they gave an average of 8.8 points under the condition that all the metadata are given and not annotated by themselves.

Overall, none of the subjects had serious problems grasping the cognitive concept of semantic regions and using the MediaFinder interface. They liked the concepts of spatial organization and dynamic grouping of the personal media data based on the different models.

4.2 Evaluation of Semantic Regions Construction

This second 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 mental 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 both how to use semantic regions and how to construct them.

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 minutes. 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 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 t heir pictures as the background images of the created regions) And see how many pi ctures were taken for each person between 1995 and 2000. (Specify Conjunctive Se mantics)

- 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 the subjects were able to complete all 4 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, both novice and expert users could create their own models with MediaFinder after just about 30 minutes of use. In addition, it was interesting to observe that no subject created a region that had more than two conjunctive semantics.

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 to 10) for MediaFinder interface, they marked an average of 7.8 points, which was not much different but 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 questionnaires 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 MediaFinder Interface

The two usability studies demonstrate that users, with some training, are able to construct their own mental models for personal media management with MediaFinder. The studies also helped to identify improvements to the user interface that reduce the need for semantic specification and provide a control panel for semantics.

The MediaFinder user interface has been redesigned similarly to the Microsoft File Explorer so that users can feel familiar with the interface in specifying the semantics for a region and annotating the 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 steps for region construction and semantics specification.

Users can choose a category tab and select a category tag from the semantic control panel. The selected category tag can be dragged onto the MediaFinder main panel. When the dragged tag is dropped on the main panel, MediaFinder automatically creates 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 mental models.

5. Conclusion

Semantic Regions is a conceptual model, implemented in the MediaFinder user interface that allows users to rapidly explore and manage personal media items stored on their personal computers. Users can dynamically construct a variety of mental 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.

Usability studies on the MediaFinder interface revealed benefits, cognitive issues, and usability concerns. Both novice and experienced users succeeded in grasping the cognitive concept of semantic regions and using the MediaFinder interface. Users have a preference for the concepts of spatial organization of information and dynamic regrouping of personal media items based on their mental models and specific tasks. This research on semantic regions makes five distinct contributions:

- **Conceptual model:** *Semantic Regions* provides a model of spatial organization and dynamic reorganization of personal media data based on users' mental 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' mental 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.

- **Implementation:** The MediaFinder uses *Semantic Regions*, the *fling-and-flock* metaphor, and flexible metadata categorization capabilities to support exploration and management of personal media data sets. Implemented in Visual Basic .Net, MediaFinder uses object-oriented design techniques that support the use of subclassing to easily add new classes of features.
- **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.

Acknowledgements:

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

References

Adar, E., Kargar, D., Stein, L.A., Haystack: per-user information environments. *In Proceedings of the eighth international conference on Information and knowledge management*, 413 - 422, ACM Press, 1999.

Ahlberg, C., Shneiderman B., Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays, *Proceedings of the ACM CHI '94 Conference*, 313-317, 1994.

Ahlberg, C., Williamson, C., Shneiderman, B., Dynamic Queries for Information Exploration: An Implementation and Evaluation, *Proceedings of the ACM CHI '92 Conference*, 619-626, 1992.

Bell, C. G., The Cyber All Project: A Personal Store for Everything, *Microsoft Research Technical Report MSR-2000-75*, July, 2000.

Bell, G., A Personal Digital Store, *Communications of the ACM, Vol. 44, No. 1* (January 2001), 86 - 91, 2001.

Bell, G., Gray, J., Digital Immortality, *Communications of the ACM, Vol. 44, No. 3* (March 2001), 28 – 30, 2001.

Card, S., Mackinlay, J., Shneiderman B., Readings in Information Visualization: Using Vision to Think, *Morgan Kaufmann*, 1999.

Chen, H., Dumais, S.T., Bringing order to the web: Automatically categorizing search results, *Proceedings of the ACM CHI'00 conference,* 145-152, 2000.

Chuah, M.C., Dynamic aggregation with circular visual designs, *Proceedings of the IEEE Symposium on Information Visualization*, 35 - 43, October, 1998.

Cockbrun A., Mckenzie B., Evaluating the Effectiveness of Spatial Memory in 2D and 3D Physical and Virtual Environments, *Proceedings of the ACM CHI2002 conference*, 203 - 210, 2002.

Dourish, P., Edwards, W. K., LaMarca, A., Salisbury, M., Presto: An Experimental Architecture for Fluid Interactive Document Spaces, *ACM Transactions on Computer-Human Interaction, 6(2),* 1999.

Fertig, S., Freeman, E., Gelernter, D., Lifestreams: An Alternative to the Desktop Metaphor, *Proceedings of the ACM CHI96 Conference*, New York, 410 – 411, 1996.

Gifford, D.K., Jouvelot, P., Sheldon, M.A., O'Toole, Jr., J.W., Semantic File Systems, *13th ACM Symposium on Operating Systems Principles*, October, 1991.

Goldstein, J., Roth, S. F., Using Aggregation and Dynamic Queries for Exploring Large Data Sets, *Proceedings of the ACM CHI '94 Conference*, 23-29, 1994.

Greene, S., Tanin, E., Plaisant, C., Shneiderman, B., Mushlin, R., Olsen, L., Major, G., Johns, S, The End of Zero-Hit Queries: Query Previews for NASA's Global Change Master Directory, *International Journal of Digital Libraries*, Vol. 2, No. 2 (3), 79 - 90, 1999.

Krasner, G.E., Pope, S.T., A Cookbook for Using The Model-View-Controller User Interface Paradigm in Smalltalk-80, *Journal of Object-Oriented Programming*, 1(3), 26 - 49, 1988.

Lamming, M., Forget-me-not: intimate computing in support of human memory, *FRIEND21, '94, International Symposium on Next Generation Human Interface*, Meguro Gajoen, Japan, 1994.

Long, B., TimeStore: Exploring Time-Based Filing. *Unpublished paper*, University of Toronto, 1994.

Mander, R., Salomon, G., Wong, Y.Y., A pile metaphor for supporting casual organization of information. *Proceedings of the ACM CHI92 Conference*, 627-634. 1992.

Mayo, R.N., Reprint of the Factoid Webpage, *Technical Note TN-60*, Compaq Computer Corporation, Western Research Laboratory, July, 2001.

Moore, G.E., Gigabits and Gigabucks, *University Video Corporation Distinguished Lecture,* 1996.

Nardi, B.A., Whittaker, S., Isaacs , E., Creech, M., Johnson, J., Hainsworth, J., Integrating communication and information through ContactMap, *Communications of the ACM*, Volume 45, Number 4, pp. 89-95, 2002.

North, C., Shneiderman, B., Component-Based, User-Constructed, Multiple-View Visualization, *Extended Abstracts of CHI '2001*, 201-202, 2001.

Pepys Project at XRCE, http://www.xrce.xerox.com/programs/mds/past-projects/pepys.html

Plaisant, C., Milash, B., Rose, A., Widoff, S., Shneiderman, B., LifeLines: Visualizing Personal Histories. *Proceedings of the ACM CHI96 Conference*, 221-227, 1996.

Quan, D., Bakshi, K., Huynh, D., and Karger, D. R., User interfaces for supporting multiple categorization. *In Proceedings of Interact 2003,* 2003.

Rekimoto, J., TimeScape: A Time Machine for the Desktop Environment, *Proceedings of the ACM CHI '99 Conference*, 1999.

Robertson, G., Czerwinski, M., Larson, K., Robbins, D. C., Thiel, D., van Dantzich, M., Data Mountain: Using Spatial Memory for Document Management, *Proceedings of ACM UIST98*, 153-162, November, 1998.

Roth, S., Lucas, P., Senn, J., Gomberg, C., Burks, M., Stroffolino, P., Kolojejchick, J., Dunmire, C., Visage: A User Interface Environment for Exploring Information, *Proceedings of Information Visualization*, IEEE, 3 - 12, 1996.

Roth, S.F., Kolojejchick, J., Mattis, J., Chuah, M., SageTools: An Intelligent Environment for Sketching, Browsing, and Customizing Data Graphics, *Proceedings of the ACM CHI95 Conference*, 409 - 410, 1995.

Shipman, F., Airhart, R., Hsieh, H., Maloor, P., Moore, J. M., and Shah, D., Visual and spatial communication and task organization using the visual knowledge builder. *In Proceedings of the 2001 International ACM SIGGROUP Conference on Supporting GroupWork*, 260-269, 2001.

Shneiderman, B., Inventing Discovery Tools: Combining Information Visualization with Data Mining, *Proceedings of Discovery Science 2001*, Springer-Verlag, Washington DC, 2001.

Shneiderman, B., Kang, H., Direct Annotation: A Drag-and-Drop Strategy for Labeling Photos, *Proceedings of International Conference Information Visualisation* (IV2000), 2000.

Tang, L., Dynamic Aggregation, CMSC838 class project in spring 2001, http://www.cs.umd.edu/class/spring2001/cmsc838b/Project/Lida_Tang/index.html

The Diary Agent, *Project Web Page*, http://www.d.kth.se/~d93-jlu/Mobile/diaryagent.html

Williamson, C., Shneiderman, B., The Dynamic Home Finder: Evaluating Dynamic Queries in a Real-Estate Information Exploration System, *Proceedings of ACM SIGIR '92 Conference*, 338 - 346, 1992.

Zellweger, P., Mackinlay, J., Good, L., Stefik, M., and Baudisch, P., City Lights: Contextual Views in Minimal Space. *Extended Abstract of CHI2003 ACM Conference on Human Factors in Computing Systems,* 2001.

Rao, R., Card, S.K., The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus+Context Visualization for Tabular Information, *Proceedings of the ACM CHI94 Conference on Human Factors in Computing Systems*, 318-322, April, 1994.