Implicit Brushing and Target Snapping: Data Exploration and Sense-making on Large Displays

Xiaohua Sun[*], Patrick Chiu, Jeffrey Huang[†], Maribeth Back, Wolf Polak

FX Palo Alto Laboratory
3400 Hillview Ave., Bldg. 4
Palo Alto, CA 94304, USA
{lastname}@fxpal.com

[*]MIT
Department of Architecture
Cambridge, MA 02139, USA
xhsun@mit.edu

[†]Harvard University
Graduate School of Design
Cambridge, MA 02138, USA
jhuang@gsd.harvard.edu

ABSTRACT
During grouping tasks for data exploration and sense-making, the criteria are normally not well-defined. When users are bringing together data objects thought to be similar in some way, implicit brushing continually detects for groups on the freeform workspace, analyzes the groups' text content or metadata, and draws attention to related data by displaying visual hints and animation. This provides helpful tips for further grouping, group meaning refinement and structure discovery. The sense-making process is further enhanced by retrieving relevant information from a database or network during the brushing. Closely related to implicit brushing, target snapping provides a useful means to move a data object to one of its related groups on a large display. Natural dynamics and smooth animations also help to prevent distractions and allow users to concentrate on the grouping and thinking tasks. Two different prototype applications, note grouping for brainstorming and photo browsing, demonstrate the general applicability of the technique.

Categories and Subject Descriptors
H.5.2 [**Information Interfaces and Presentation (e.g., HCI)**]: User Interfaces - Graphical user interfaces (GUI).

General Terms
Design, Human Factors.

Keywords
Visual interfaces, information visualization, sense-making, grouping, large displays

1. INTRODUCTION
Grouping and labeling are basic tasks that support data analysis, exploration, and sense-making. In the physical world, a simple yet effective way of working makes use of work surfaces on walls and tables with data objects in the form of note cards or Post-it notes. By putting notes similar in some way into spatially separate groups and labeling them, users can explore various interpretations of the data, see them at a higher level as labeled sets, and gain a better understanding of the data. Established work practices that employ this kind of highly visual interaction include methodologies such as contextual design [3], and management and decision-making processes like the KJ Method [9] which is well-known in Japan. Even for more common activities like organizing and browsing media collections (e.g. photos), grouping and labeling by themes is often performed.

AVI '06, May 23–26, 2006, Venezia, Italy.

When electronic surfaces such as tabletop and wall displays are used as the workspace for the grouping and labeling, the data exploration and sense-making process can be facilitated by the data analysis and retrieval capability of computers. Designing the visualization and interaction techniques that employ these advantages to an appropriate extent can be challenging. When the data has ambiguities and uncertainties – as real data tend to have – the computerized display should not go too far in automatically clustering the data; it must give users an opportunity to interpret and direct the data around the display. In a multi-user situation where discussion and negotiation take place, the system must not be intrusive. From these considerations, the familiar thematic maps driven by force diagrams (e.g. [7]) can be overly active.

The approach presented in this paper is to visualize the grouping potential of data objects but not move them prematurely. The grouping potential is displayed through group color decoration and arrows pointing to related groups. The grouping potential is based on analysis of the groups' text and metadata. Relevant objects from a database or network can also be retrieved into the foreground of the workspace with their grouping potential shown. Users can analyze the data aided by these visual hints, decide how to group the data and how to label the groups. The visual hints and animated data retrieval can save the users much effort in finding candidate group members. Looking at the multifaceted meaning of potential group members may also help clarify the meaning of the group or even redirect the purpose of the grouping.

The brushing – finding of related data and generating visual hints – can sometimes be initialized by the text explicitly stated on the group labels. In practice, however, the labels and descriptive phrases that formalize the groups are often not settled until the later stages of the sense-making process; during the early stages the meaning of the groups are ambiguous and negotiable.

The *implicit brushing* technique that we designed continually does group detection and analyzes group member text and metadata to initiate the brushing implicitly based on this dynamic group information without requiring the user to formally label the groups. On freeform workspaces, detecting implicit structures is an approach that has been shown to support effective interaction (e.g. [6], [11]). Our system detects groups based on overlap, and for each group the aggregate information with score of each term weighted by relevance and number of occurrences provides an objective description of the group that covers all related aspects. Both this analysis and the further utilization of the result in brushing would be overwhelming for the user to perform manually without computational aid.

When doing brushing, visual hints are generated for the related data; some of this data is already on the workspace while other data is retrieved. Hence, our brushing process not only draws attention to data objects on the workspace, but also brings into the workspace related data from a database or network. This feature employs relevance feedback by generating a query for each group, boosting the common terms that occur in a group. While the retrieval mechanism is similar to existing systems (e.g. [1], [5]), our interface is designed to support this in a highly visual manner as the graphical objects on the workspace are manipulated during exploration and sense-making.

In our visualization, the Z dimension (or scale change), is used to represent this process of information retrieval triggered by the grouping manipulations. Depending on an object's relevance status, it is positioned at different depths and appears in different sizes. See Fig. 3 and 5. Four layers are employed for: (1) working set, (2) relevant peripheral data, (3) previously retrieved peripheral data, (4) peripheral dataset. The peripheral dataset is represented by tiny dots in the background using different transparency to show their age. Displaying the tiny dots in Layer 4 is optional; e.g. it is feasible to do this for peripheral data coming from a relatively small well-defined collection in a database but not for data coming from the entire Web. The retrieved peripheral objects move in and out along the depth as their relevance status is changed during the grouping and labeling performed by the users. At the same time as retrieving relevant data, the relevance status of previously brushed data is checked. The visual hints on data objects no longer relevant will be removed; those retrieved from the peripheral database will be sent back to the "previously retrieved peripheral data" layer.

Considering the size of the typical displays our interface is intended for, the *target snapping* technique is developed to provide convenience in moving data objects to their potential groups. Standard drag-and-drop operations from desktop PCs do not work well on a large display. Thus, various techniques have been developed (e.g. [2], [4], [8]). While it is possible to incorporate some of these techniques, our design provides a simple solution: make the arrow visual hints from the brushing interactive. By tapping on one of its arrowheads, a data object is sent across the display with animation to the relevant group.

In the next section we explain and illustrate the visualization, interaction, and retrieval mechanisms in more detail.

2. VISUALIZATION & INTERACTION

Designed as a virtual version of a physical workspace for data exploration through grouping, the interface we created keeps the natural dynamics of interaction as in the physical world. It has a freeform workspace with data objects scattered around and piles of blank group labels placed at the corners. See Fig. 2 and 3. Following the same convention as in reality, a group is created when a user places an object so that it overlaps another. A subtle point is that groups must be intentionally created by a user; objects that inadvertently overlap during the initial placement or retrieval by the system are not considered as group members (only the borders of group members are highlighted).

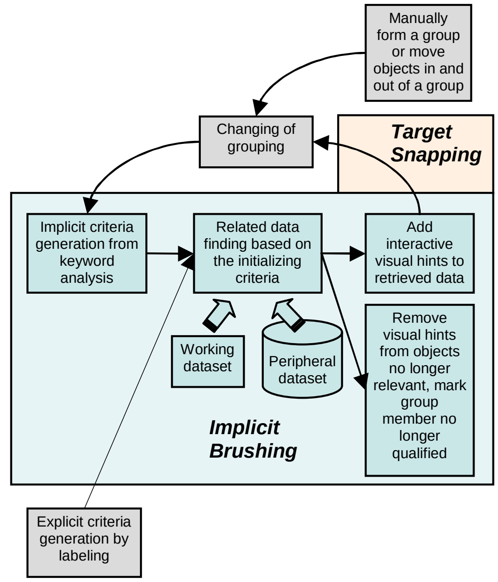

Figure 1. Flow of the dynamics

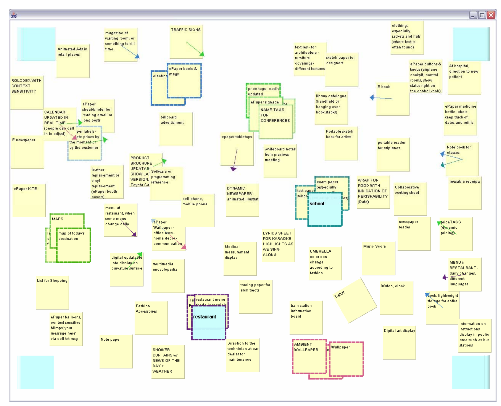

Figure 2. Screen shot of brainstorming application

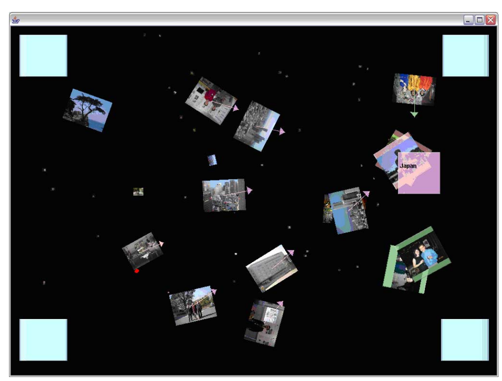

Figure 3. Screen shot photo browsing application

Figure 4. Data object examples

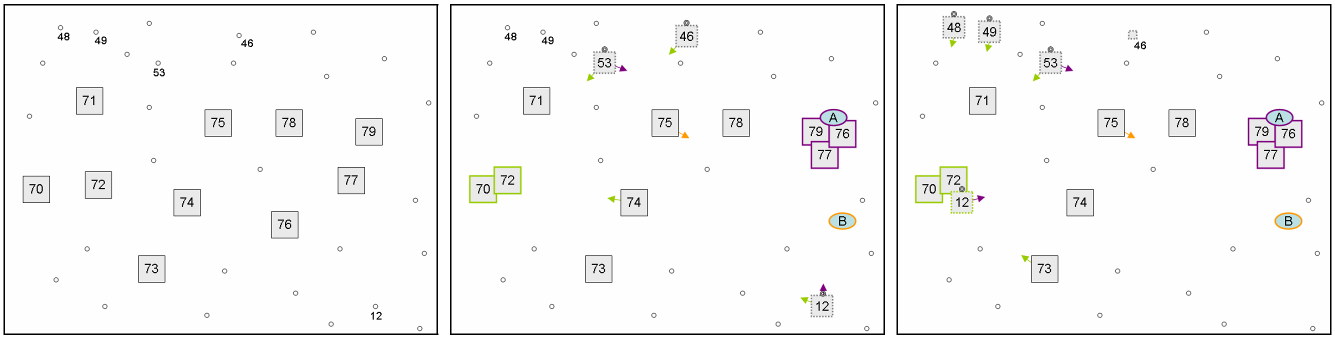

Figure 5. Schematic illustration of grouping aided by implicit brushing and target snapping

A grouping action triggers the implicit brushing, which in turn provides visual hints for further grouping and structure discovery. The flow of the dynamics is illustrated by the diagram in Fig. 1. Fig. 2 and Fig. 3 show screen shots of the interface in two prototype applications.

The data objects are typically thumbnail-sized as shown in Fig. 4. When they form or join a group, their borders are painted with the group's color, which is assigned by the system at the time of group formation. Transparency is used to show how relevant an object is to the group. By assumption, each object has text content or text metadata associated to it.

Users can tap or click on an object and drag it around. An object can be rotated; supporting orientation is important for table displays (e.g. [10]). For wall displays, rotating an object provides a way of distinguishing it. Pressing down on an object magnifies it for better viewing.

Peripheral data objects retrieved from a database or network are smaller than the objects in the working set. They are also painted a different style of borders. In addition, there is a circular widget for pinning it to the "relevant peripheral data" layer. For example, see object 53 in Fig. 5-2 and 5-3.

A label object is used in two situations: first, when the meaning of a group is getting clearer, the user may want to label it with a certain phrase; second, when the user wants to form a group with predetermined keywords, she can initialize an explicit brushing by placing a label on an empty area in the working space. In both cases, right-clicking on the label would bring up a list of phrases suggested by the system based on the text analysis of the group or the working set. When the application exits, phrases on the labels can be saved out and added to the metadata of the group members.

Dragging a label moves the whole group of overlapping member objects. Dragging a data object moves only that object regardless of whether it is part of a group or not. The rationale is that a group formalized with a label should be manipulable as a whole unit since a user has intended for it to be a unit, whereas a group that is not labeled should be more amenable to membership changes.

During a session, the working set is shown in the foreground layer, and the peripheral data from a database shown in the background. Initially, the objects are randomly scattered on the display. Schematically, this is illustrated in Fig. 5-1. A grouping and brushing example will be explained in detail aided by illustrations in Fig. 5.

When the user moves object 72 together with 70 with overlap, a group is formed, the group is detected and implicit brushing is triggered. As a result, to object 74 in the working set is added an arrow pointing to the group with the group's color (green). Objects 12, 46 and 53 from the database are retrieved to the "relevant peripheral data" layer and added arrows pointing to the group. Clicking the green arrow on object 12 moves it to group {70, 72} by target snapping. This changes the target group to {12, 70, 72}, which triggers another round of implicit brushing. As shown in Fig. 5-3, objects 48 and 49 retrieved from the database plus object 73 from the working set are now brushed as relevant. At the same time, the relevance status of the previously related data is checked. Visual hints are removed from objects 46 and 74 which are no longer relevant, and object 46 from the database is sent backward to the "previously retrieved peripheral data" layer.

In Fig. 5, the label objects are represented by ovals. The label object A has been added to the group {76, 77, 79} to formalize it. Label object B on the workspace shows an example of explicitly initializing a brushing by placing a label in an empty area; its related object 75 is brushed.

Animations are used at many places in the visualization and interaction to provide natural and vivid visual effects. When a group is moving around, all the arrows pointing to this group are continuously reoriented to keep pointing to it. Arrows on group members pointing to their potential groups are also reoriented accordingly. When arrows are added to relevant objects or removed from no longer relevant objects, blinking fade-in and blinking fade-out animations are employed to attract the attention of users. Target snapping motion is animated to make it easier for users to follow the action. Animations along the Z dimension are also employed when retrieving objects from the "peripheral data" layer to the "relevant peripheral data" layer and when sending them back to the "previously retrieved peripheral data" layer.

3. APPLICATION SCENARIOS

We prototyped the technique and informally tested it in two different application scenarios with real data. In this section we describe these scenarios and explain how implicit brushing and target snapping can support them on large displays. One scenario is for a wall display and the other for a table display (see Fig. 6).

The first application is sense-making through note grouping and labeling. For data, the KJ Method [9] was used to coordinate a brainstorming session. The data was generated using physical Post-it notes and later the sense-making was performed on an electronic display. Five people participated in an hour-long session: four were researchers on the project plus a staff member from Japan who was familiar with the method. The topic was "electronic paper applications." The session generated 72 notes; this data was collected for input into the application. In this scenario, no peripheral data, such as data from previous brainstorming sessions, was generated.

The grouping and sense-making of the collected data using our interface is carried out on a wall display system consisting of three side-by-side rear projectors with total resolution 3072x768. See Fig. 6-1. The text on note objects, each at size 125x125, is legible from a reasonable distance. The implicit brushing technique provides great help in spotting grouping candidates from the "ocean" of data objects. The target snapping is obviously superior to dragging as a way to move an object a long distance across the large display.

The text contents rendered directly on the notes of the already formed group members and the candidate objects help users get a more and more precise understanding of the various possible meanings of their related groups. For example, meaningful groups for the "electronic paper applications" brainstorming data might emerge and be labeled as "eBook", "news", "reference", etc. This grouping and labeling facilitate the process of discovering higher level categories and structures from brainstorming data.

Figure 6. Prototype application running on three-screen wall display and tabletop display

The second application is a photo browsing scenario on a table display. We built a prototype coffee table system with a 1280x768 resolution display and put this in our lab's lounge. See Fig. 6-2. A relatively small set of photos was culled by selecting highlights from business trips, company social events, along with some personal pictures. These photos were manually tagged with metadata consisting of keywords that describe the event, such as the location and the people.

The idea is that a coffee table display can function like photo albums found on coffee tables in lounges or living rooms. Visitors can browse the photos and learn more about an organization or a family. With our setup, initially a featured event (e.g. most recent) is shown on the table. The user may add or remove events. The photos from these chosen events comprise the working set and are shown in the foreground; photos from the rest of the events are shown in the background as peripheral data.

The implicit brushing and target snapping technique facilitates rapid exploration and different ways of looking at the collection. For example, a user might put into a pile photos from a conference in Japan due to her initial interest in that conference. Besides photos of this specific event, data retrieved during the implicit brushing might include other photos of Japan or photos of the same conference at other places in other years. If many of the photos in that pile contain a person "Jim", other photos of "Jim" might also be retrieved as relevant. These candidates in different directions suggest the next possible exploration direction based on the estimated potential interest of the user from photos she has grouped. Such an interaction between the user and the system makes the browsing process natural and smooth.

The interplay of the dynamics is highly compelling in this application. While the user explores and shifts her focus on the various groups and individual photos in the XY-plane of the workspace, the retrieved photos from the database flow back and forth along the Z axis depending on their relevance to the current interest of the user.

4. CONCLUSION & FUTURE WORK

In summary, this paper presented the technique of implicit brushing and target snapping for grouping data on large displays. To demonstrate the general applicability of this technique, we prototyped two quite different scenarios for sense-making with notes and exploration of photo collections. For future work, we plan to do more extensive user testing, add support for easier data input for notes and photos, and assess how well it can visualize larger datasets that arise through accumulation from multiple sessions or through a search engine to the Web.

5. REFERENCES

[1] Baldonado, M., Winograd, T. SenseMaker: an information-exploration interface supporting the contextual evolution of a user's interests. *Proceedings of CHI '97*, pp. 11-18.

[2] Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P. Bederson, B., Zierlinger, A. Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-operated Systems. *Proceedings of Interact '03*, pp. 57-64.

[3] Beyer, B., Holtzblatt, K. *Contextual Design: Defining Customer-Centered Systems*, Morgan Kaufmann, 1997.

[4] Collomb, M., Hascoet, M., Baudisch, P., Lee, B. Improving drag-and-drop on wall-size displays. *Proceedings of GI '05*, pp 25-32.

[5] Cutting, D., Karger, D., Pedersen, J., Tukey, J. Scatter/Gather: a cluster-based approach to browsing large document collections. *Proc. of SIGIR '92*, pp. 318-329.

[6] Moran, T., Chiu, P., van Melle, W., Kurtenbach, G. Implicit structures for pen-based systems within a freeform interaction paradigm. *Proceedings of CHI '95*, pp. 487-494.

[7] Olsen, K., Korfhage, R., Sochats, K., Spring, M. Williams, J. Visualization of a Document Collection: the VIBE System, *Information Processing & Management*, 29(1): 69-82 (1993).

[8] Rekimoto, J. Pick-and-Drop: A direct manipulation technique for multiple computer environments. *Proceedings of UIST '97*, ACM Press, pp. 31-39.

[9] Scupin, R. The KJ Method: A Technique for Analyzing Data Derived from Japanese Ethnology. *Human Organization*, 56(2): 233-237 (1997).

[10] Shen, C., Everitt, K., Ryall, K. UbiTable: Impromptu face-to-face collaboration on horizontal interactive surfaces. *Proceedings of UbiComp'03*, pp. 281-288.

[11] Shipman, F., Marshall, C., Moran, T. Finding and using implicit structure in human-organized spatial information layouts. *Proceedings of CHI '95*, pp. 346-353.