

# TeamSearch: Comparing Techniques for Co-Present Collaborative Search of Digital Media

Meredith Ringel Morris, Andreas Paepcke, Terry Winograd  
Stanford University  
{merrie, paepcke, winograd}@cs.stanford.edu

## Abstract

*Interactive tables can enhance small-group co-located collaborative work in many domains. One application enabled by this new technology is co-present, collaborative search for digital content. For example, a group of students could sit around an interactive table and search for digital images to use in a report. We have developed TeamSearch, an application that enables this type of activity by supporting group specification of Boolean-style queries. We explore whether TeamSearch should consider all group members' activities as contributing to a single query or should interpret them as separate, parallel search requests. The results reveal that both strategies are similarly efficient, but that collective query formation has advantages in terms of enhancing group collaboration and awareness, allowing users to bootstrap query-specification skills, and personal preference. This suggests that team-centric UIs may offer benefits beyond the "staples" of efficiency and result quality that are usually considered when designing search interfaces.*

## 1. Introduction

Small-group activity is common in many aspects of daily life, including work, learning, and recreation. Single display groupware (SDG) refers to a class of computing technology that supports work among small, co-present groups around one shared display [18]. The display provides group members with a shared context and focus of attention, conversation, and activity, and is a common paradigm for supporting co-located cooperative work.

The recent introduction of computationally-enhanced tables offers software designers the opportunity to develop applications that support co-located collaboration among groups of users, such as collaborative exploration of digital libraries by co-present groups. Collaborative search for information is not well-supported by current SDG technology, although data exploration is a task often accomplished

in small-group settings. For example, a group of business colleagues or students might search through a repository of charts and documents to compile relevant bits of information into a report or presentation, or a family might search through a collection of personal digital photographs to assemble a themed album. We have developed TeamSearch, an application supporting co-present collaborative search of metadata-tagged digital content on an interactive table (Figure 1).



**Figure 1. A four-person group uses TeamSearch at a DiamondTouch table to find photos from a metadata-tagged repository.**

## 2. Related Work

Several systems enable collaborative work around interactive tables. The UbiTable [13] allows two users to transfer digital media from their laptops to a tabletop display where it can be shared and annotated. RoomPlanner [24] allows users to create furniture arrangements using special gestures on an interactive tabletop. The InteracTable [19] allows groups to annotate digital content on a computationally-enhanced table. ConnecTables [21] allow users to combine mobile desks to create a larger horizontal work surface and share and exchange documents. SoundTracker [8] is a tabletop application for group music exploration. These projects all address the creation and/or manipulation of digital content using tables; in contrast, TeamSearch explores using interactive tables to support collaborative search of digital content.

The Personal Digital Historian (PDH) [12] is a tabletop application that supports storytelling by allowing a group of users to collectively interact with a set of digital photos. PDH allows users to query the photo collection along one of four possible dimensions – who is in a photo, what event is depicted, or where or when the photo was taken. However, the hardware used in the PDH system supports interaction by only one group member at a time, so PDH’s creators were unable to explore truly collaborative query formation.

Many research and commercial systems, such as PhotoFinder [16], Fotofile [5], Adobe’s Photoshop Album<sup>1</sup>, Apple’s iPhoto<sup>2</sup>, and Google’s Picasa<sup>3</sup>, offer photo tagging and searching capabilities. These systems are all designed for operation by a single user, while TeamSearch focuses on multi-user, collaborative search of digital content.


Studies have shown that many people have trouble specifying Boolean queries [3, 20]. To make query formulation more accessible, systems such as [9, 11, 15, 23] allow a single user to specify Boolean queries using a visual or tactile scheme rather than an abstract language. TeamSearch extends the concept of visual query formation to include collaborative queries. Prior work on collaborative information retrieval, such as the Ariadne system [22] focuses on allowing remote users to assist each other, while our focus is on co-located collaborative search. However, the focus of this paper is on exploring different styles of collaborative query formation rather than on contributing a novel style of non-verbal query specification.

### 3. The TeamSearch System

TeamSearch is a multi-user application that allows four-member groups to collaboratively search collections of digital content, such as photos, that have been previously associated with relevant metadata. Users form Boolean-style<sup>4</sup> queries by arranging circular “query tokens” on the tabletop (Figure 2).

TeamSearch users sit around an 85.6 cm x 64.2 cm DiamondTouch table (Figure 1). The table is run by a 3.2 GHz PC, and a display is top-projected onto the table’s surface with an SXGA projector. The DiamondTouch [1] is a touch-sensitive surface that is

capable of receiving simultaneous touch input from each of up to four users, and of identifying the (x,y) location of each touch on its surface as well as identifying which of the four users are associated with each touch event. The TeamSearch software is written using the DiamondSpin tabletop interface toolkit [14].



**Figure 2. The starting configuration of TeamSearch consists of several components: (a) The collection of photos being searched is represented as a pile in the center of the table. (b) The shaded rectangular regions on each side of the table are where thumbnails that match the current query will be displayed. (c) A pile of *query tokens* (round objects labeled “?”) is located on each side of the table. (d) Circular widgets represent the schema of the photo collection’s metadata. Each circle corresponds to a category (e.g., “people” or “location”), and each wedge within a circle corresponds to a specific metadata value for that category (e.g., “Alex,” “Larry,” or “Lisa”). Users search the photo collection by placing query tokens on top of target metadata values, and thumbnails of the matching images are shown in the shaded rectangular regions. Touching a thumbnail brings the corresponding photo to the top of the pile so users can inspect and interact with it.<sup>5</sup>**

<sup>1</sup> <http://www.adobe.com/products/photoshopalbum>

<sup>2</sup> <http://www.apple.com/ilife/iphoto>

<sup>3</sup> <http://www.picasa.com>

<sup>4</sup> Since our goal is to explore support for co-located collaborative querying and not to contribute to the literature on visual query languages, TeamSearch does not offer complete Boolean expressivity, but rather interprets all token combinations as an “AND” (during pilot testing we found that this simplification made it easier for users to specify queries, which was not surprising given prior studies on the difficulty many people have with the Boolean conceptual model, such as [3, 20]).

<sup>5</sup> Note that this screenshot has been modified – the sizes of the tokens, photos, and circular widgets have been enlarged relative to the size of the table in order to enhance legibility for publication. The inset depicts the actual relative scales of the interface components, and correctly shows the substantial amount of open space available on the table both for manipulating photographs as well as potentially displaying additional metadata widgets. Note that figures 3 and 4 have also been edited in this manner to enhance legibility.

When TeamSearch is initialized, all of the photos in the current repository appear in a virtual pile in the table's center (Figure 2a). These photos have previously been manually tagged with several categories of metadata (some metadata is also automatically added, using techniques described in [10]). A rectangular area in front of each of the four users is initially blank – this is the area where query results, shown as thumbnails corresponding to query-satisfying images, will be shown (Figure 2b). To each user's left is a circular token marked with a “?” – this is a query token (Figure 2c). A user can move a query token by touching and dragging it about the surface of the table with his fingertip. When a token is moved from its original location, a new one appears underneath it – essentially, there is an infinite pile of query tokens for each user. Near the center of the table are several circular widgets, which are subdivided into wedges. Each circle represents a category of metadata (e.g., “location”), and each wedge within that circle is labeled with a specific possible value for that category (e.g., “Italy,” “Israel,” “Sri Lanka”) (Figure 2d).<sup>6</sup>


First, we explain how a single user creates a query with TeamSearch. We then describe how groups can collaboratively query the photo repository.

Suppose User X wants to find all of the photos in the collection that were taken in Sri Lanka. Manually searching through the pile would be inefficient since there are several hundred photos, and even if he looked through them, he might not remember which ones were from Sri Lanka. Instead, User X decides to query the collection. He takes one of the query tokens from his token pile and drags it with his finger into the wedge marked “Sri Lanka” within the circular widget that contains the “location” metadata category. He places the token on that wedge and releases it. In the shaded rectangular region in front of User X, several thumbnail images appear. Each of these thumbnails corresponds to an image from the collection that satisfies the criterion “location=Sri Lanka”. In order to find the original, full-resolution image, User X can press on one of the thumbnails with his finger. The corresponding image will move up to the top of the pile in the center of the table and will blink to aid User X in locating it. User X can then touch that image with his finger and move it around the table, resize or reorient it, view other metadata associated with it, etc. Suppose User X wants to further revise his query to

<sup>6</sup> Available screen space limits the total number of metadata categories/values that can be simultaneously displayed. TeamSearch could be adapted for use with large schemata using several techniques, such as shrinking infrequently-used widgets or organizing metadata hierarchically and displaying one level at a time. Detailed discussion of scaling techniques is beyond the scope of this paper.

find a more specific image – he wants to find an image from Sri Lanka that has his brother Larry in it. To refine his query, he takes another token from his token pile, and places this one on the wedge marked “Larry” within the circular widget representing the “people” category. The display of thumbnails in front of him updates to show matches only for photos satisfying the query “location=Sri Lanka AND people=Larry.”

This querying technique can be extended in order to permit all four people sitting around the table to work collaboratively on a search task. We consider two implementation alternatives that offer different interpretations of how the system should process simultaneous token placements by members of the group – *collective* and *parallel* querying.



**Figure 3. TeamSearch with *collective query tokens*: all tokens contribute to a single query.**



**Figure 4. TeamSearch with *parallel query tokens*: each group member's tokens (distinguished by color) form distinct queries.**

Under the *collective query tokens* implementation, when tokens are placed onto the circular widgets the system interprets all tokens collectively as a single query no matter which group member placed them. For example, if User X placed a token on “Larry” and User Y placed a token on “Lisa” and a token on “Sri Lanka,” then the result would be a single query “location=Sri Lanka AND people=Larry AND people=Lisa,” and the thumbnails that matched that query would be displayed in front of each user (Figure 3).

*Parallel query tokens* offer a more relaxed interpretation of collaborative querying, which permits individual group members to form distinct queries in parallel with other users at the table. This design is influenced by observations of group work indicating that small-group tasks tend to transition between periods of tightly-coupled group activity interspersed with periods of more loosely-coupled individual work [2, 7].

Under this implementation, when tokens are placed onto circular widgets the system interprets all tokens placed by each individual user as a single query, for a maximum of four queries at any one time (one per user). Each user’s query tokens are a different color, to make this distinction clear. Using parallel query tokens, if User X placed a token on “Larry” and User Y placed a token on “Lisa” and a token on “Sri Lanka,” then the result would be that the thumbnails matching the query “people=Larry” would be shown in front of User X, the thumbnails matching “people=Lisa AND location=Sri Lanka” would be shown in front of User Y, and no thumbnails at all would be shown in front of the two users who placed no tokens (Figure 4).

When developing TeamSearch, it was not apparent whether the collective or parallel query scheme was more appropriate for use by co-located groups collaboratively searching through digital collections towards a common goal. Prior work on the tradeoffs between group-oriented versus individual-oriented designs for CSCW systems have focused on distributed systems [4, 17], but have not explored how these issues apply to SDG. To better understand the benefits and drawbacks of each querying style, we conducted an empirical study. The purpose of this experiment was to clarify questions relevant to designing interface mechanisms to support co-located collaborative search, such as: (1) Does either design allow people to reach their search goals more effectively? (2) Does either design facilitate more efficient searching? (3) Does either design promote more effective collaboration among group members? (4) Will users have strong preferences for either of the designs?

## 4. Evaluation

We recruited sixteen paid subjects to participate in our study. Subjects’ ages ranged from twenty to thirty years old, and they were evenly split between genders. Participants completed the experiment in groups of four users at a time, for a total of four groups. The experiment had a within-groups design, with each group completing two search tasks using two different sets of photos with analogous metadata schemata, with one task using collective query tokens and one using parallel tokens. The order of photo sets and token types was balanced using a Latin Square design.

In each condition, a collection of seventy-five digital images was shown on the table. Each image in the set was associated with four categories of metadata: people, location, event, and year. There were five possible values for each of the four categories (e.g., year={2000 | 2001 | 2002 | 2003 | 2004}). A single photo could have multiple people associated with it, but only a maximum of one value each for the other three categories. The photos were not from the subjects’ personal collections, so they had to rely on querying, rather than recognition or brute force search, to find specific photos. Groups were told to choose a subset of the images for the purpose of making hard-copy prints to place in a photo album. The requirement for the album was that each person, location, event type, and year must be represented in at least one photo. A single photo could satisfy multiple requirements simultaneously. Groups were encouraged to find a minimal set of photos that satisfied the requirements for their album in order to lower printing costs.

When a group was satisfied that the set of photos they had chosen for printing covered all of the required values and was minimal, they told the experimenter that they were finished. They were then given a questionnaire to complete individually, asking them to evaluate certain aspects of their experience. The same procedure was then repeated using the other token style and a new set of photos.

Throughout the study, all user interactions with the table were logged by our software (e.g., movements of query tokens, interactions with photos and thumbnails, etc.).

## 5. Results

The results from our evaluation of TeamSearch can be grouped by four themes: the quality of the answers found; the efficiency of each search technique; the impact of each interface on group collaboration; and user preference data.

## 5.1 Quality

We use two measures to gauge the quality of the outcome. First, did the chosen set of photos provide complete coverage of each of the twenty metadata values (four categories x five values each)? In each condition every group achieved full coverage, so there was no difference between the two techniques with regard to this aspect of quality. The second quality measure regards the size of the chosen set of photos. According to the instructions given to each group, answers that were as close as possible to the minimal number of necessary photos were desirable. Groups were not told what this number was. Although all groups did not select the optimal set of photos in all conditions (an optimal answer could contain 5 photos), the average size of the final set did not differ significantly regardless of token type: the mean final set size was 6.5 photos with the collective tokens and 7.25 photos with the personal tokens, which is not a statistically significant difference ( $t(3)=1.19$ ,  $p=.32$ ). Thus, both interfaces were similar in terms of quality of the outcome of the search task.

## 5.2 Efficiency

Several measures of efficiency can be used to analyze the two query-token schemes. First, we can look at the total task time in each condition. The mean time with the collective tokens was 12.65 minutes, while it was 11.50 minutes with the parallel tokens. This difference is statistically indistinguishable ( $t(3)=.50$ ,  $p=.65$ ). For all groups, whichever condition they experienced second was faster (an average of 10.09 minutes compared to 14.06 minutes in the first session), reflecting a reliable learning effect in terms of more efficient use of TeamSearch ( $t(3)=5.90$ ,  $p<.01$ ). Groups experienced a larger learning effect (5.11 minute time decrease vs. 2.82 minute time decrease) when they worked first with collective tokens followed by parallel tokens rather than vice-versa ( $t(3)=4.85$ ,  $p<.02$ ). We conjecture that this effect could be due to users who had more difficulty understanding how to make Boolean queries learning from teammates during the early exposure to the closely-coupled collective token interface. These users were then better prepared to work more independently with the parallel query tokens.

Another measure of efficiency is to look at the rate of querying (i.e., total number of queries made / total time). This measure reveals a significant difference between the techniques, with collective tokens yielding a rate of .056 queries/sec, while the parallel tokens yielded a higher rate of .110 queries/sec ( $t(3)=4.56$ ,

$p<.02$ ). By the query-rate standard, the parallel tokens resulted in the ability to form queries more quickly.

Another perspective on the efficiency issue is to explore not how many queries were made, but how sophisticated each query was. For example, a single complex query might have the expressive power of two simpler queries. In this light, the more complex query could be viewed as a more efficient method of answering a question. We examined whether either of the two implementations of TeamSearch encouraged the formation of more sophisticated queries by measuring the most complex query (in terms of number of tokens combined into a single query) formed by each group in each condition. Groups were able to achieve similarly complex queries with each interface (an avg. max. complexity of 5 tokens with the collective interface and of 3.81 tokens with the parallel interface), ( $t(3)=1.34$ ,  $p=.27$ ), so neither technique had an efficiency advantage with respect to this criterion.

## 5.3 Collaboration

One important aspect of an interface for co-located group search is that it facilitates collaboration among group members. There are several metrics we can explore to examine the impact that each interface design had on groups' collaborative activities.

Examining the balance of work among group members is a key aspect of evaluating the system's impact on collaboration. A group with a very skewed balance of work (e.g., all queries contributed by only one of the four group members) can be considered to be collaborating less than a group where all members contributed more equally to the task. We can examine the interaction logs to see how many queries were contributed by each user within a group, and then calculate the standard deviation for each group of the number of queries contributed (i.e., number of tokens placed on metadata values) by each member. This allows us to summarize how balanced the group's participation was in contributing queries (i.e., a smaller standard deviation within a group indicates more balanced participation) (note that this measure does not take verbal contributions into account). Taking the mean of this per-group standard deviation score across each of the groups within each condition, we find the mean is 5.78 with the collective tokens and 9.09 with the parallel tokens ( $t(3)=4.89$ ,  $p<.02$ ), indicating a more balanced distribution of query formation among group members when using the collective query token interface.

Awareness of other group members' activities is another important aspect of collaboration, particularly if the search activity is intended as part of an educational goal, since higher awareness of other

group members' actions could result in more incidental learning [6]. We measured awareness by having participants make three judgments on the questionnaire they were given immediately following each experimental condition. Subjects were asked to indicate the number of queries they thought they had personally executed during the activity, the combined total number of queries they thought all four group members had executed, and how many members of the group (from 0 to 3) they felt had executed more queries than they had personally. We compared these assessments to the actual data recorded by our system to check accuracy. More accurate assessments of these values would indicate higher awareness of one's own and/or others' interactions with TeamSearch. The mean difference between the perceived and actual number of queries done personally by each group member was 5.84 with collective tokens and 11.25 with parallel tokens ( $t(15)=2.95$ ,  $p<.01$ ). The mean difference between the perceived and actual number of queries done by all group members was 20.38 with collective tokens and 35.53 with parallel tokens ( $t(15)=2.54$ ,  $p<.03$ ). The mean difference between the perceived and actual number of group members who had contributed more queries than the survey respondent was .81 with collective tokens and 1.19 with parallel tokens ( $t(15)=2.42$ ,  $p<.03$ ). In all three of these cases, the lower mean difference for collective tokens indicates a higher awareness than with parallel tokens.

**Table 1. Collective tokens received higher mean ratings on a 7-point Likert scale regarding their impact on collaboration.**

	Collective tokens	Parallel tokens	p-value
I worked closely with the other members of my group to accomplish this task.	5.75	4.88	$p<.04$
Members of the group communicated with each other effectively.	5.75	5	$p\leq.05$
The group worked effectively as a team on this task.	5.75	4.81	$p<.03$

We also gathered participants' subjective self-reports regarding various aspects of collaboration. These self-report data indicate that the collective query tokens facilitated more effective collaboration with group members than did the parallel query tokens. Subjects answered three Likert-scale questions (7-point

scale) relating to various aspects of collaboration. For each of the three questions (Table 1), the average rating was significantly better for the collective tokens.

## 5.4 Satisfaction

After completing both conditions, each participant individually completed a questionnaire asking her to make comparisons between the two conditions. On this survey, the majority of subjects (10 of 16, 62.5%) reported a preference for the collective interface as compared to the parallel interface. Subjects also reported greater satisfaction with the task outcome when using the collective tokens (as indicated by mean scores given on a 7-point Likert scale for: "I am satisfied with the set of photos that my group selected"), with a mean of 6.0 for the collective tokens and 4.88 for the parallel tokens ( $t(15)=3.74$ ,  $p<.01$ ).

## 6. Discussion

Based on the quantitative and qualitative data gathered during our study, we can revisit the design questions that initially motivated our exploration of the comparative strengths and weaknesses of the collective and parallel query token interfaces for co-located collaborative search of digital photo collections. The increased awareness, more equitable distribution of work, and heightened satisfaction with the collective tokens suggests that the more team-centric interface offers benefits beyond the "staples" of efficiency and result quality that are usually considered when designing interfaces for searching digital media.

*Does either design allow people to reach their search goals more effectively?* Groups were able to achieve their search goals for the study task (complete coverage of all categories/values, and small answer-set size) equally well with either search interface.

*Does either design facilitate more efficient searching?* We had initially expected that the parallel query tokens might facilitate more efficient searching, since they provide group members with more independence and flexibility, allowing the group to present several queries to the system simultaneously (up to one query per user). However, we found only minimal efficiency benefits to the parallel scheme, which resulted in a faster query formation rate than the collective interface, but which did not significantly impact total time spent on the search task or query complexity. Based on the results of our study, it seems that the potential efficiency benefits introduced by parallelism might have been cancelled out by the learning benefits of the collective tokens, which seem to have helped "weaker" group members more quickly

catch on to how to use TeamSearch by providing the opportunity for them to work in synchrony with more query-savvy group members. It is likely that, with longer-term use, the efficiency benefits of the parallel scheme would become more pronounced; however, our results regarding collaboration and satisfaction suggest that some of the less tangible benefits of camaraderie and teamwork might still bend preferences toward the collective query interface.

*Does either design promote more effective collaboration among group members?* Because the collective tokens facilitate a more closely-coupled work style, we suspected that they would result in an increased sense of collaboration among users. This suspicion was borne out by subjects' self-reports of several dimensions of collaborative activity. Feelings of working closely as a team and of communicating well with the group were rated significantly higher with the collective interface.

The collective interface also resulted in higher awareness by participants about both their own and other group members' contributions to the task. While we had expected that the collective interface would facilitate more awareness about others' contributions, we had thought that the parallel tokens might facilitate increased *self-awareness* by more explicitly highlighting individual contributions (through the color-coding of the tokens). One possible explanation for the increased personal awareness in the collective condition is that people felt more of a need to recall and emphasize their own contribution in this case, since the collective interface did not make it obvious who had contributed which parts of the queries.

Finally, the collective interface resulted in more even distribution of the work of query formation among group members. Again, we were initially surprised by this result, since prior work [8] found that adding more individual flexibility to a group tabletop system resulted in more equal distribution of work among the group members; for that reason, we had expected that the parallel query tokens might result in more balanced participation, while the use of the collective tokens might end up being dominated by a single, aggressive group member. However, the challenging nature of forming Boolean-style queries [3, 20] might have been a key factor in changing the nature of participation in this task (as compared to the task studied in [8] which was a tabletop entertainment application rather than a tabletop search application). With the parallel interface, participants who were more confused by query formation might have felt unable to contribute a query on their own, but with the collective tokens often the more dominant individuals would direct other group members where to place tokens in order to help the group form a collective query, thus

encouraging participation from all group members. The increased confidence of "weaker" users with the collective tokens is reflected by the questionnaire responses of the only two participants in our study who had never heard of the concept of Boolean queries. These two subjects indicated more agreement with the statement "I was confused about how to form queries" for the parallel tokens interface (rating of 4 and 5 on a 7-point scale) than with the same statement about the collective tokens (rating of 2 and 3).

*Will users have strong preferences for either of the designs?* Although subjects ranked both interfaces as similarly easy to use and understand, the majority of participants in our study preferred using the collective, rather than the parallel, tokens, and also reported greater satisfaction with the final set of photos their group selected with the collective interface. Perception of teamwork was highly correlated with self-reported satisfaction with the outcome ( $r=.525$ ,  $p<.04$ ). We were surprised by this, since our ongoing work on collaborative photo-labeling has found that users prefer individual sets of controls when performing labeling tasks on an interactive table. Perhaps the more challenging nature of the search task as compared to the labeling task influenced the preference for more closely-coupled teamwork in this situation.

## 7. Conclusion

We have introduced TeamSearch, a tabletop application that enables small, co-located groups to search for digital photos from a metadata-tagged repository. Because co-located group query formation is a relatively unexplored domain, we needed to answer basic questions to improve the design of the TeamSearch interface – whether an interface for group query formation should consider search constraints provided by each group member as contributing to a single, complex query (collective query token interface) or whether each group member's searches should be executed individually (parallel query token interface). Our evaluation found only minor differences between the two interfaces in terms of search quality and efficiency, but found that the collective interface offered significant benefits in terms of facilitating stronger collaboration and awareness among group members and in terms of users' preferences. The advantages of the collective interface may be related to the difficulties of forming Boolean-style queries and the fact that this interface allows group members with weaker query-formation skills to learn from other group members. Our evaluation of these two alternative querying interfaces for TeamSearch is a valuable first step toward understanding the unique

requirements for designing successful tabletop interfaces that enable co-located groups to access digital media repositories.

## 8. Acknowledgements

We thank MERL for donating a DiamondTouch table, and the AT&T Labs Fellowship Program.

## 9. References

- [1] Dietz, P., and Leigh, D. DiamondTouch: A Multi-User Touch Technology. *UIST 2001*, 219-226.
- [2] Elwart-Keys, M., Halonen, D., Horton, M., Kass, R., and Scott, P. User Interface Requirements for Face to Face Groupware. *CHI 1990*, 295-301.
- [3] Green, S.L., Devlin, S.J., Cannata, P.E., and Gomez, L.M. No IFs, ANDs, or ORs: A Study of Database Querying. *International Journal of Man-Machine Studies*, 32, 3 (1990), 303-326.
- [4] Gutwin, C. and Greenberg, S. Design for Individuals, Design for Groups: Tradeoffs Between Power and Workspace Awareness. *CSCW 1998*, 207 – 216.
- [5] Kuchinsky, A., Pering, C., Creech, M.L., Freeze, D., Serra, B., and Gwizdka, J. Fotofile: A Consumer Multimedia Organization and Retrieval System. *CHI 1999*, 496-503.
- [6] Lave, J. and Wenger, E. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge, UK, 1991.
- [7] Mandviwalla, M. and Olfman, L. What Do Groups Need? A Proposed Set of Generic Groupware Requirements. *Transactions on Computer-Human Interaction*, 1(3), 1994, 245-268.
- [8] Morris, M.R., Morris, D., and Winograd, T. Individual Audio Channels with Single Display Groupware: Effects on Communication and Task Strategy. *CSCW 2004*, 242-251.
- [9] Murray, N., Paton, N., and Goble, C. Kaleidoquery: A Visual Query Language for Object Databases. *AVI 1998*, 247-257.
- [10] Naaman, M., Harada, S., Wang, Q., Garcia-Molina, H., and Paepcke, A. Context Data in Geo-Referenced Digital Photo Collections. *ACM Multimedia 2004*, 196-203.
- [11] Pane, J. and Myers, B. Improving User Performance on Boolean Queries. *CHI 2000 Extended Abstracts*, 269-270.
- [12] Shen, C., Lesh, N., Vernier, F., Forlines, C., and Frost, J. Sharing and Building Digital group Histories. *CSCW 2002*, 324-333.
- [13] Shen, C., Everitt, K., and Ryall, K. UbiTable: Impromptu Face-to-Face Collaboration on Horizontal Interactive Surfaces. *UbiComp 2003*, 281-288.
- [14] Shen, C., Vernier, F., Forlines, C., and Ringel, M. DiamondSpin: An Extensible Toolkit for Around-the-Table Interaction. *CHI 2004*, 167-174.
- [15] Shin, D., and Nori, R. CBM: A Visual Query Interface Model Based on Annotated Cartoon Diagrams. *AVI 2000*, 242-245.
- [16] Shneiderman, B. and Kang, H. Direct Annotation: A Drag-and-Drop Strategy for Labeling Photos. *International Conference on Information Visualization*, May 2000.
- [17] Stefik, M., Bobrow, D.G., Foster, G., Lanning, S., and Tatar, D. WYSIWIS Revised: Early Experiences with Multiuser Interfaces. *ACM Transactions on Office Information Systems*, 5(2), April 1987, 147-167.
- [18] Stewart, J., Bederson, B., and Druin, A. Single Display Groupware: A Model for Co-present Collaboration. *CHI 1999*, 286-293.
- [19] Streitz, N., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R. i-LAND: An Interactive Landscape for Creativity and Innovation. *CHI 1999*, 120-127.
- [20] Tanaka, J. The Perfect Search. *Newsweek*, 134(13), September 27, 1999, p. 71.
- [21] Tandler, P., Prante, T., Müller-Tomfelde, C., Streitz, N., and Steinmetz, R. ConnecTables: Dynamic Coupling of Displays for the Flexible Creation of Shared Workspaces. *UIST 2001*, 11-20.
- [22] Twidale, M. and Nichols, D. Designing Interfaces to Support Collaboration in Information Retrieval. *Interacting with Computers*, 10(2), 1998, 177-93.
- [23] Ullmer, B., Ishii, H., and Jacob, R. Tangible Query Interfaces: Physically Constrained Tokens for Manipulating Database Queries. *Interact 2003*, 279-286.
- [24] Wu, M. and Balakrishnan, R. Multi-Finger and Whole Hand Gestural Interaction Techniques for Multi-User Tabletop Displays. *UIST 2003*, 193-202.