Contents lists available at ScienceDirect

Information Processing and Management

Search on surfaces: Exploring the potential of interactive tabletops for collaborative search tasks

Meredith Ringel Morris [a,*], Danyel Fisher [a], Daniel Wigdor [b]

[a] *Microsoft Research, Redmond, WA, USA*
[b] *Microsoft Surface, Redmond, WA, USA*

ARTICLE INFO

ABSTRACT

Collaborative information seeking often takes place in co-located settings; such opportunities may be planned (business colleagues meeting in a conference room or students working together in a library) or spontaneous (family members gathered in their living room or friends meeting at a café). Surface computing technologies (i.e., interactive tabletops) hold great potential for enhancing collaborative information seeking activities. Such devices provide engaging direct manipulation interactions, facilitate awareness of collaborators' activities, and afford spatial organization of content. However, current tabletop technologies also present several challenges that creators of collaborative information seeking system must account for in their designs. In this article, we explore the design space for collaborative search systems on interactive tabletops, discussing the benefits and challenges of creating search applications for these devices. We discuss how features of our tabletop search prototypes TeamSearch, FourBySix Search, Cambiera, and WeSearch, illustrate different aspects of this design space.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

In the movie *The Day the Earth Stood Still* (20th Century Fox, 2008), a group of scientists gather around a tabletop computer, pointing and arguing with each other – the tabletop supports their discussion, providing additional data and detail as they work. This scenario is, of course, science fiction: few computer systems today support smooth transitions between collaborative work and data exploration.

Many types of information work are well-suited to collaboration: a team can provide a variety of perspectives, bring a spectrum of expertise, and parallelize tasks. Hence, group work is an important component of many types of information work (Fidel et al., 2000; Hansen & Jarvelin, 2005; Morris, 2008) and educational activities (Amershi & Morris, 2008; Large, Beheshti, & Rahman, 2002; Talja, 2002; Twidale, Nichols, & Paice, 1997). When using printed media, workgroups often collaborate on information retrieval tasks (Fidel et al., 2000), and traditional horizontal surfaces (desks and tables) are often the location where such media are examined, discussed, and made sense of. Yet when groups move to digital media, that comfortable collaboration is lost: users are forced to gather around a single PC (perhaps projected on a meeting-room wall for communal viewing), or to work separately at their own computers.

Recently, researchers have begun creating digital systems that support collaborative information retrieval (Amershi & Morris, 2008; Morris & Horvitz, 2007; Pickens, Golovchinsky, Shah, Qvarfordt, & Back, 2008). Researchers have also made much progress in designing digital desks and tabletops (Dietz & Leigh, 2001; Haller et al., 2006; Han, 2005; Hartmann,

---

* Corresponding author.
  E-mail addresses: merrie@microsoft.com (M.R. Morris), danyelf@microsoft.com (D. Fisher), dwigdor@microsoft.com (D. Wigdor).

Morris, Benko, & Wilson, 2009; Johanson, Fox, & Winograd, 2002) to support 21st-century work and educational practices. The affordances of digital tabletops offer great promise for facilitating collaborative search experiences; indeed, we have begun to explore some of these synergies in our own work (Hartmann et al., 2009; Isenberg & Fisher, 2009; Morris, Paepcke, & Winograd, 2006; Morris, Lombardo, & Wigdor, 2010), which we discuss in more detail below.

In this article, we explore the potential of interactive tables for supporting collaborative search activities. We first discuss the unique affordances of tables, and consider both the benefits that these offer to collaborative search tasks, as well as the challenges they present; we then discuss the design space of table + search applications. We describe four systems from our own research, TeamSearch, FourBySix Search, Cambiera, and WeSearch, and reflect on how they fit into this design space. We close by discussing promising directions for future work at the intersection of these two fields.

1.1. Related work: systems for co-located collaborative search

A few researchers have begun to explore collaborative search on a tabletop form-factor. Físchlár-DT (Smeaton, Lee, Foley, & McGivney, 2006) is a tabletop video search system that enables pairs of users to work together when searching through a video archive on a DiamondTouch table (Fig. 1). The system was inspired by (and uses data from) the TRECVid video retrieval challenge. On Físchlár-DT, video search is conducted by typing keywords on soft keyboards. Matching keyframes can then be examined using direct-touch interactions, played on an auxiliary monitor, or saved by being dragged to a special region of the table. Keyframes themselves can also be used to initiate a search, through a "find similar" action. The designers of Físchlár-DT were particularly interested in the tradeoffs in designing a tabletop search system for maximum awareness among partners versus designing for maximum efficiency (i.e., supporting highly parallel work styles), and found that higher-awareness versions of their application were more preferred by end-users.

The Personal Digital Historian (PDH) Shen, Lesh, Vernier, Forlines, and Frost, (2002) is a tabletop application designed for collaborative browsing of photo collections, in support of storytelling activities (Fig. 2). While designed more for browsing than for search, the PDH provides four views that query the tagged photo collection – "who" (filtering the collection based on who is in a photo), "what" (filtering based on what event is depicted), "where" (filtering based on the photo's originating location), and "when" (filtering based on timestamp). The PDH tabletop hardware supported only a single input at a time, so collaborative discussion was supported, but not collaborative input to the system itself.

Other than Físchlár-DT, PDH, and our own work on the TeamSearch, FourBySix Search, Cambiera, and WeSearch systems (which we discuss in Section 3), we are not aware of any other research efforts that explore the potential of interactive tabletops for collaborative information seeking tasks. However, a closely-related body of work is the exploration of interfaces for co-located collaborative search using other, non-tabletop, technologies, such as multiple co-located PCs, PDAs, and/or mobile phones. The tasks that motivate many of these co-located search systems (education, information work, reviewing photos and video, etc.) could potentially be rich areas for support by tabletop systems, as well. The face-to-face working style enabled by a tabletop display has the potential to support many of the co-located collaborative search tasks addressed by these systems, while providing higher awareness levels than when each user is focused on her own personal device.
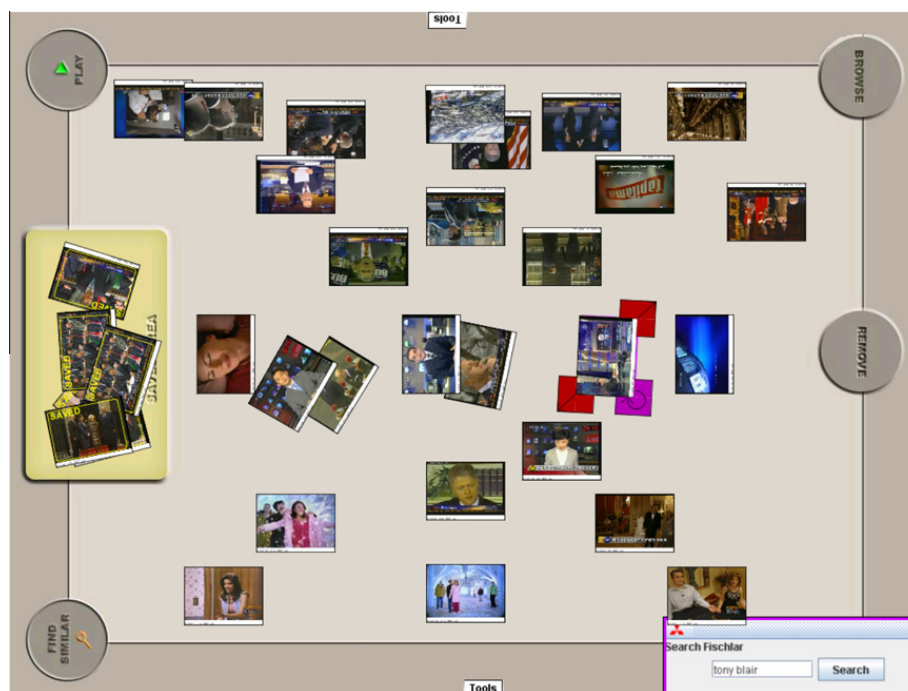

Fig. 1. The Físchlár-DT system enabled partners to search through a video library. (Image courtesy of Alan F. Smeaton.)

Fig. 2. The Personal Digital Historian supported group storytelling by enabling filtering of photo collections along several dimensions. (Image courtesy of Chia Shen.)

Like Físchlár-DT and PDH, the Cerchiamo system (Pickens et al., 2008) is also focused on multi-media search, supporting search over a video archive. Cerchiamo is designed to support two co-located searchers, sitting side-by-side with their own PCs, augmented by a large, shared wall display. Each partner's personal PC displays information relevant to his distinct role in the search process (either the "prospector", who enters query terms, or the "miner", who triages results), while the shared display shows data relevant to both parties, i.e., the overall progress of the search task.

In addition to co-located multi-media search support, some researchers have been exploring tools for co-located collaborative search of Web pages. Such systems tend to have a separate device (such as a phone or PDA) for each group member, sometimes augmented by a shared PC or wall-mounted display. CoSearch (Amershi & Morris, 2008) combines mobile phones (one per group member) with a shared display to support collaborative search. The phones' keypads can be used to enter queries and the phones' joysticks can be used to select links to follow; queuing mechanisms in the browser on the shared display help coordinate the input from each group member's phone. WebGlance (Paek et al., 2004) focuses on collaborative Web browsing, rather than search *per se*, enabling users to control a browser on a large, shared display by issuing commands from their individual PDAs. Maekawa, Hara, and Nishio's (2006) system divides a Web page into several non-overlapping segments, and shows each segment on a different group member's phone, in order to enable the group to divide and conquer the task of visual search within a page. The Query By Argument system (Blackwell, Stringer, Toye, & Rode, 2004) is a rhetorical argument construction system that provides groups of students with RFID-tagged tangible objects representing key concepts from items in a document collection; manipulating these tangibles, such as by grouping them together in specific ways, results in re-calculating the relevance rankings of items in the document set, thus making the students aware of additional information that may be related to their current argument.

2. Designing collaborative search applications for tabletops: benefits and challenges

The affordances of horizontal computing surfaces are in many respects harmonious with the needs of collaborative search systems; however, there are still several challenges in adapting this emerging computing form-factor to information retrieval tasks. In this section, we discuss both the benefits and challenges of using tabletops as a platform for collaborative search.

2.1. Benefits of tabletops for collaborative search

The face-to-face working style supported by interactive tabletops tends to encourage more equitable work styles than the shoulder-to-shoulder work style supported by other types of single display groupware, such as interactive walls (Rogers & Lindley, 2004). Interactive tabletops support co-located collaboration, but can also support remote or mixed-presence collaboration in situations where multiple tabletops are distributed at different locations; virtual arrangements of seating patterns (Tuddenham & Robinson, 2009) and virtual embodiments such as video arms (Izadi et al., 2007) can help preserve the engagement of a face-to-face around-the-table experience even for distributed groups.

Awareness of collaborators' searching activities has been found to be an important aspect of collaborative search systems (Morris & Horvitz, 2007; Paul & Morris, 2009; Smeaton et al., 2006). Although awareness features can be designed into such systems, collaborators sharing a single display report higher awareness and better group work experiences than those with separate displays, even when all the displays are side-by-side (Amershi & Morris, 2008). The large, shared display provided by an interactive tabletop, then, can offer benefit for collaborative search tasks, by providing shared focus and awareness among group members.

Large, horizontal surfaces afford spreading, piling, and organizing content. Such affordances are well-suited to sensemaking (Russell, Stefik, Pirolli, & Card, 1993), which is an important part of the collaborative search process (Paul & Morris, 2009) since search in itself is never a group's end-goal; rather, groups engage in collaborative search tasks in order to achieve some larger goal, such as to prepare a report or presentation, plan an itinerary, or make a purchasing decision (Morris, 2008). Hence, the organizational affordances of tabletops offer great potential for rich collaborative sensemaking applications.

And last, but not least, horizontal displays can support placement of tangible props, and can respond and react to them using technologies such as recognition of RFID or optical tags (Hartmann et al., 2009; Wilson & Sarin, 2007). Future technologies may eventually have high-fidelity vision systems that enable OCR of un-tagged documents. Such capabilities can enable groups to incorporate traditional media into their search process. For example, a page from a book, a photograph, or a physical object could be used as input to initiate a search task.

2.2. Challenges of tabletops for collaborative search

While we believe that the aforementioned affordances of tables make them well-suited to collaborative search tasks, it is important to acknowledge that there are some challenges to designing successful tabletop search systems. Interactive tabletops are still an emerging technology, and some aspects of their design are still being explored by the tabletop research community.

One of the most engaging aspects of interactive tabletops are their ability to provide direct-touch manipulations (Benko, Morris, Brush, & Wilson, 2009). However, direct-touch interaction is inherently less precise than interactions on traditional PCs; pointing with fingers is less exact than mouse-based pointing (stylus-based tabletops address this concern to some extent (Haller et al., 2006)), and typing on soft keyboards is much slower than on physical ones (Hartmann et al., 2009). The tedium of text entry on tabletops using soft keyboards (Benko et al., 2009; Hinrichs, Hancock, Collins, & Carpendale, 2007) can be particularly challenging for search applications, where text entry is often a key part of the activity.

Although current tabletop technologies tend to be physically larger than PC displays (the Microsoft Surface, for example, measures 30″ diagonally, and the DiamondTouch table measures 42″ diagonally), the resolution of these displays is often not much greater than the resolution of a standard monitor (and sometimes even less), since XGA (1024×768 pixels) is still the most common and affordable projector resolution, and most tabletop technologies are projection-based. This results in spatially large, but low-resolution displays. Even as tabletop hardware progresses to include larger, higher-resolution surfaces, space limitations will continue to be a challenge. The clutter caused by displaying content for an entire group of people on one display is one of the fundamental challenges of single display groupware systems (Stewart, Bederson, & Druin, 1999), and exploratory search tasks test the limits of this issue since they often involve a large amount of content (Paul & Morris, 2009).

Orientation is a perennial challenge for tabletop displays (Shen, Vernier, Forlines, & Ringel, 2004), since content that is right-side-up for one user is upside-down for anyone seated across the table. Some automated solutions are possible, including lazy-susan-style UIs (Shen et al., 2004), replication of content for all users (Morris, Paepcke, Winograd, & Stamberger, 2006), "magnetization" of the UI in one direction at a time (Shen et al., 2004), and affordances for manually re-orienting content (Shen et al., 2004). Orientation issues become particularly critical during collaborative search tasks, as search often involves dealing with large amounts of textual data, and dense text is particularly orientation-sensitive (Wigdor & Balakrishnan, 2005); some domains, such as image or video search, are somewhat more resistant to this issue, although even in these cases considering orientation is still an important aspect of tabletop search application design, due to the important communicative properties conveyed by orientation in group settings (Kruger, Carpendale, Scott, & Greenberg, 2003).

As we have described, current tabletop technologies' limitations regarding text entry, clutter, and orientation can make it particularly challenging to adapt these technologies to productivity tasks (Benko et al., 2009; Morris, Brush, & Meyers, 2008); however, with continued advances in hardware and software from the tabletop community, as well as careful application design by the search community, these challenges can be overcome. In the next section, we describe four applications that illustrate several approaches to addressing these issues.

3. Exploring the design space

Tabletop search applications can span a large design space that encompasses many group work styles and many types of search tasks. Here, we articulate what we see as the key dimensions of this design space based on our own projects and on others' related work. By articulating this space, we hope to better understand how existing work in this area (such as our own systems, discussed below) complement each other, and what areas remain unexplored; such unexplored areas of the space suggest interesting directions for future research, as we discuss in Section 4.

As with all collaborative search systems, it is important both to consider the mechanics of the collaboration itself, as well as the specifics of the type of search activity being supported. Thus, we subdivide our design space into the "group configuration" and "search task" sub-sections. Table 1 illustrates the salient dimensions of the design space.

The first aspect of group configuration that we consider is the collaborative style the system is meant to support – very tightly-coupled collaborations, loosely-coupled parallel work, or styles that support transitions among both (an articulation of these different styles can be found in Tang, Tory, Po, Neumann, and Carpendale (2006)). Group size is also a relevant factor; we must consider whether a system is meant to support searches among pairs of users, small groups (3–6 users), or large groups (i.e., all of a meeting's attendees, or an entire classroom full of students). The physical arrangement of these collaborators is also an important design choice: they can be distributed in space, each at their own tabletop; co-located around a single surface; or in a mixed arrangement, with some team members co-located and others remote. Finally, understanding the larger ecology of the group's collaboration is vital – are they conducting their entire search task on the tabletop, or do they need to import data they gathered individually before the collaborative phase began? Do they need to export the data for consumption by themselves or others after the tabletop session has ended? Do they have auxiliary inputs (paper documents, laptops, phones, etc.) or outputs (such as vertical displays for detailed reading) that they will use in conjunction with the surface?

Several aspects of the search task also bear consideration. First, the *type* of search task must be considered. A search application can support Web search, multi-media search, search over group members' own document collections, or search over a specialized database, such as medical databases or other types of digital libraries. These differences are important both in considering the size of expected result sets and the method of presenting individual results. How the search will be conducted is also a key consideration – will users specify inputs via keywords, via a visual query language, via touch and gesture interactions, or via tangible props (such as pages from a book, physical icons, or other artifacts), or will the inputs be suggested to the user via automatic, "intelligent" means, such as based on the user's or group's recent actions? It is also important to consider what aspects of the search task the system will address: the application may support the *process* of search, *understanding the products* of search (i.e., sensemaking), or preparing a final "report" containing the groups' findings (or perhaps some combination of several of these phases).

Our own research has explored some of the ways in which interactive tabletops can support collaborative search tasks. In the next sections, we describe four prototype applications, TeamSearch, FourBySix Search, Cambiera, and WeSearch, and reflect on how they address different aspects of this design space and how they resolve the challenges of adapting tabletops to search tasks.

3.1. TeamSearch

Our first prototype, TeamSearch (Morris, Paepcke, & Winograd, 2006), is designed to help small groups of co-located users search through tagged image repositories; for example, helping a family select photos from their collection in order to assemble a themed website or album, or helping students or co-workers assemble images to support preparation of a group presentation or report. Table 2 shows where the TeamSearch system falls in our design space.

TeamSearch uses a 42″ diagonal DiamondTouch (Dietz & Leigh, 2001) table, which is a top-projected, capacitive-touch technology that is able to differentiate the identities of up to four users based on their seat via capacitive coupling connecting a special seating pad, the user, and the device (Fig. 3). In this case, the projected display has SXGA (1280 × 1024) resolution. The application is designed for searching over a tagged photo collection; the photo tagging itself is not part of TeamSearch, but can occur using other single-user (e.g., flickr.com) or collaborative (e.g., Morris et al., 2006) tagging systems. The application can be configured to use a photo + tag repository of the group's choosing. Upon initialization, TeamSearch parses the tag structure and creates a circular widget representing each category of tags (i.e., people, location, etc.); each circle is subdivided into sectors representing the range of values for that tag category (i.e., for the category "people", values might include "mom", "dad", "Mary", "Sue", "Bob", etc.).

In addition to the circular tag widgets, other key components of the UI include the photos themselves, which can be moved, rotated, and resized via gestures in a fashion consistent with other multi-touch applications. The UI also includes a rectangular *tray* area located near each user's seat, which displays thumbnails of images matching the current search.

Table 1
Design space of tabletop collaborative search applications.

Group configuration					
Collaborative style	Tightly-coupled	Transitioning	Parallel work		
Group size	Pairs	Small teams	Large groups		
Location	Co-located	Distributed	Mixed-presence		
Device ecology	Table	Table + additional outputs/displays	Table + additional input devices/objects	Import data	Export data
Search task					
Application domain	Web	Multi-media	Personal documents	Database	
Search inputs	Keywords	Visual language	Touch/gestures	Tangible	A.I.
Analysis stages	Search	Sensemaking	Reporting		

Table 2
Coverage of the table + search design space by our own and others' systems.

| System	Group configuration									Device ecology					Search task											
	Collaboration style			Group size			Location			Table	Extra outputs	Extra inputs	Import	Export	Application domain				Search inputs					Analysis stage		
	Tightly-coupled	Transitioning	Parallel work	Pairs	Small teams	Large groups	Co-located	Distributed	Mixed-presence						Web	Multi-media	Personal documents	Data base	Keywords	Visual lang.	Touch/gesture	Tangible	A.I.	Search	Sense making	Reporting
TeamSearch	X		X		X		X			X						X				X				X		
FourBySix Search		X			X		X			X					X				X					X		
Cambiera		X	X	X			X			X								X	X		X			X		
WeSearch		X	X		X		X			X				X	X				X		X	X		X	X	
Fischlár-DT	X			X			X			X	X					X			X		X	X		X		
PDH	X				X		X			X						X								X		

Fig. 3. A group gathers around a DiamondTouch table to conduct a joint search through an image library using TeamSearch.

Pressing a thumbnail from the results tray brings the corresponding photo to the top of the pile of images (if it was buried beneath others) and causes it to blink, in order to help locate photos. And, finally, the TeamSearch UI includes *query tokens*, small circular widgets labeled with a "?" that can be instantiated via a context menu.

TeamSearch supports two styles of query tokens. *Collective query tokens*, rendered in black, and *parallel query tokens*, rendered in four distinct colors, one color for each group member. Groups can query the image repository by placing query tokens on top of the sectors of the circular tag widgets. All collective tokens are interpreted as a single query regardless of which user placed them (all covered sectors from one tag category are considered disjunctively, while each tag category that has any sectors covered is considered conjunctively) (Fig. 4). Thumbnails of the images matching the collective query are replicated in the results tray of all group members. Parallel tokens, on the other hand, are interpreted per-user; the results of a user's parallel query are shown in her own results tray, but not in other group members' (Fig. 5).

TeamSearch represents one approach to addressing the challenges of creating collaborative search applications for the tabletop form-factor. The text input challenge is mitigated by the media search domain – by utilizing tagging metadata, which is becoming increasingly common (human-computation games like ESP (von Ahn & Dabbish, 2004) aim to eventually provide tags for all images on the Web, for example), we create a faceted search experience that enables the use of visual tokens to specify which of the facets are of interest, thus avoiding the need to enter text. The *result trays* are designed to address the challenge of clutter (touching a search result thumbnail causes the corresponding original photo to become more salient even if it is lost in the clutter of photos and tag widgets on the tabletop). The trays also address the challenge of orientation, since in the collective mode results are replicated in each group member's result tray for easy multi-angle viewing.
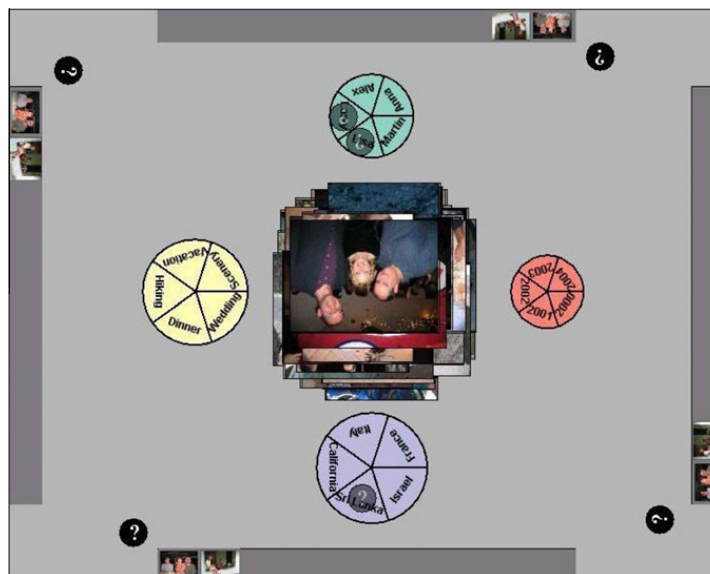
Fig. 4. The TeamSearch UI with collective query tokens. The current query specified is "(location = "Sri Lanka") and (people = "Lisa" or "Larry"). The two photos that match this query are replicated in each group members' results tray.

Fig. 5. The TeamSearch UI with parallel query tokens. The blue user and the yellow user have created two separate queries (location = "Sri Lanka" and people = "Lisa" in yellow; people = "Larry" in blue), and each sees the results for his own query in his own results tray. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
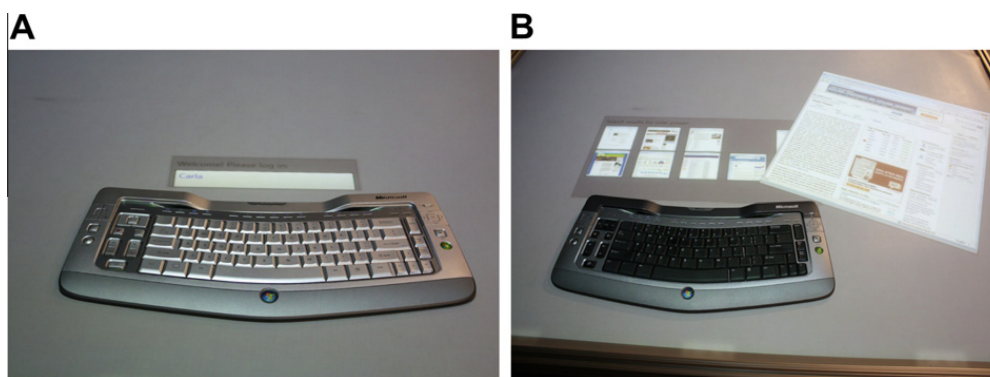
3.2. FourBySix Search

Our next prototype, FourBySix Search, supports group Web search; for example, FourBySix Search might be used by a group of students gathering material for a joint homework assignment. Table 2 illustrates the portions of the design space addressed by FourBySix Search.

The FourBySix Search application runs on a custom-built surface, named FourBySix (Hartmann et al., 2009), that is 4 feet wide × 6 feet long (Fig. 6). The system is top-projected by two tiled XGA projectors (for a total resolution of 1024 × 1536), and illuminated from beneath with infrared light; two cameras beneath the table sense IR light reflected from users' fingers, enabling multi-touch interactions. Objects that have been augmented with IR-reflecting tags can also be recognized by the system. FourBySix can recognize the location, orientation, and identity of each of several tagged keyboards placed on its surface. Simultaneous text entry from multiple keyboards is enabled through use of the Microsoft Raw Input SDK (Hartmann et al., 2009).

FourBySix Search supports transitioning between loosely-coupled and closely-coupled collaborative search styles by using the relative position and orientation of keyboards as proxies for the relative positions of the collaborators themselves. When collaborators' keyboards are above a threshold distance from each other, their inputs are interpreted separately; when a user types keywords, they are projected in a textbox displayed immediately above his keyboard; this box tracks the keyboard if it is moved to another part of the surface (Fig. 7A). When the user hits "enter", executing the search, results are displayed as thumbnails in a grid that is again co-located above the keyboard itself, in order to make it clear which set of search results was fetched by which user (Fig. 7B).

Fig. 6. A group gathers around FourBySix, a large, standing-height tabletop that supports mouse and keyboard input, to conduct a collaborative search task.

Fig. 7. (A) Placing the keyboard on FourBySix and typing creates a search box adjacent to that keyboard. (B) When keyboards are far from each other on the table's surface, each keyboard's input is treated individually, and the corresponding search results are projected adjacent to the keyboard. Touch interactions can be used to interact with an individual result in more detail.

When keyboards are placed within threshold distance of each other, the application transitions to a more tightly-coupled work style, in which the keywords entered by each of the adjacent keyboards are jointly interpreted as part of a single query. The nature of this interpretation is parameterized by the relative orientation of the keyboards. If keyboards are joined in an open, side-by-side configuration (Fig. 8A), the keywords are combined disjunctively, so as to "open" (broaden) the search. On the other hand, if the keyboards are joined in a closed, face-to-face orientation (Fig. 8B), the keywords are combined conjunctively, so as to "close" (restrict) the search.

FourBySix Search deals with the challenge of text entry on tabletops by augmenting the tables with physical input devices, in order to enable fast and reliable text entry in the text-intensive domain of Web search. The system addresses the challenge of clutter by introducing a custom-built, large form-factor surface. The challenge of orientation is addressed by having search boxes and search results track the keyboard that spawned them, in order to orient content correctly for the issuing user; results can also be dragged away from a keyboard in order to be manually re-oriented for another group member using multi-touch gestures, if desired.

3.3. Cambiera

Our third prototype, Cambiera (Isenberg & Fisher, 2009), is designed as a small-group collaborative search tool for textual databases (Fig. 9). It was motivated by the visual analytics tasks performed by intelligence analysts. Such tasks involve sorting through many textual documents to uncover a hidden narrative; the data set from the IEEE VAST competition provides a way for researchers to simulate such tasks. To complete this task, users must read many documents, searching and re-searching as they uncover related patterns and themes. Cambiera is implemented on a Microsoft Surface, a 30″ diagonal, bottom-projected (1024 × 768) multi-touch surface. Table 2 shows how Cambiera fits into our design space.

Since visual analytics tasks require high awareness among group members, Cambiera is built around the concept of *collaborative brushing and linking* (CBL). Collaborative brushing and linking is named after the standard information visualization technique of brushing and linking (Buja, McDonald, Michalak, & Stuetzle, 1991), in which the same data point is linked across different visual objects, and so all instances of it are highlighted when any one of them is brushed. A similar technique highlighting related documents can be found in Sun, Chiu, Huang, Back, and Polak (2006). In Cambiera, CBL means that the
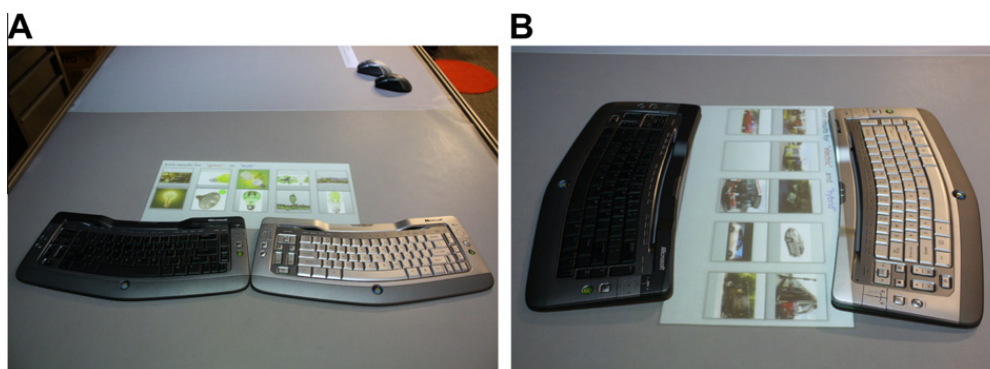
Fig. 8. When FourBySix detects that two keyboards have been moved near each other, their inputs are interpreted jointly, parameterized by their relative orientations. (A) In an open, side-by-side configuration, the search is "opened" (broadened), by displaying results for the disjunction of each keyboard's search terms. (B) In a narrow, face-to-face configuration, the search is "narrowed" (restricted), by displaying results for the conjunction of each keyboard's search terms.

Fig. 9. A pair of users collaborate around Cambiera to solve the VAST intelligence analysis challenge.
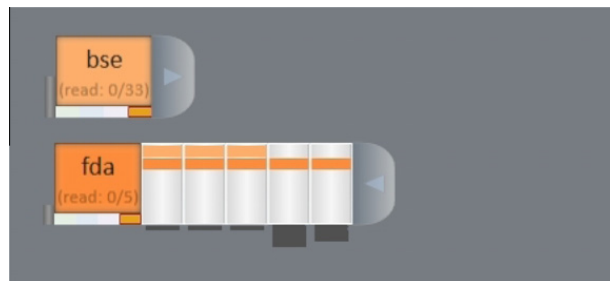

Fig. 10. Two search result boxes. The "BSE" box is in a minimized, clutter-reducing state, while the "FDA" box has been expanded to reveal collaborative brushing and linking annotations.

actions of one user are visible to other users at the places where data overlaps. For example, if the result sets for two users' queries contain some documents in common, CBL calls for a visible cue showing this overlap (Figs. 10 and 12). Similarly, if one user has read a document that is in another's results, the other user should be able to see that the document has already been read.

In Cambiera, users begin by invoking a soft keyboard, which produces a *search results box*. In (Fig. 10), a user has created two search results boxes, one searching for the term "BSE", the other for the term "FDA". The latter has been expanded to show the five hits on the term "FDA." Each of those five documents is represented by a small gray rectangle, and highlighted with a dark orange stripe, emphasizing that FDA—with its matching color—can be found in each document. In addition, three of the documents are further highlighted by a paler orange stripe, corresponding to the color of the search results box for "BSE". That pale orange tells the user that these documents can be found by both the queries "BSE" and "FDA". The saturation of the gray shading on each search result indicates how often that particular document has been read by members of the group. The user can drag out any of these rectangles from the search results box in order to create a representation of the document on-screen. Initially, such representations show the document in a small, "closed" state, in which only the title is visible, and a gesture can be used to expand such documents to the "open" state, in which the full text can be read (Fig. 11).

The critical aspect of these search boxes is that they can simultaneously display which keywords have found a given document, which user typed each of these words, and whether the document has been read. This allows the users to smoothly transition between a loosely-coupled style and a tighter one: noticing that the other user is reading a document can encourage a discussion of the document.

Cambiera alleviates the challenge of text input on tabletops by aiding users in minimizing the amount of soft-keyboard-based text input they need to do: once a user has successfully done one or more keyboard-based searches, many others can then be launched by re-using terms (via touch selections) from within the retrieved documents, allowing the user to type fairly infrequently. Additionally, persisting prior query terms through the search result box representations allows users to re-access old searches, again saving on typing. Cambiera addresses the clutter challenge by enabling data to exist at multiple scales. Search result boxes, for instance, can be minimized to show only the search term and number of results, or can be expanded to show the minimized result representations containing read-shading and CBL striping, and documents selected from search result lists can be represented in either the "closed" or "open" states. Lastly, while objects can be freely rotated, they have natural orientations; users are encouraged to sit opposite each other, and documents face that well-known
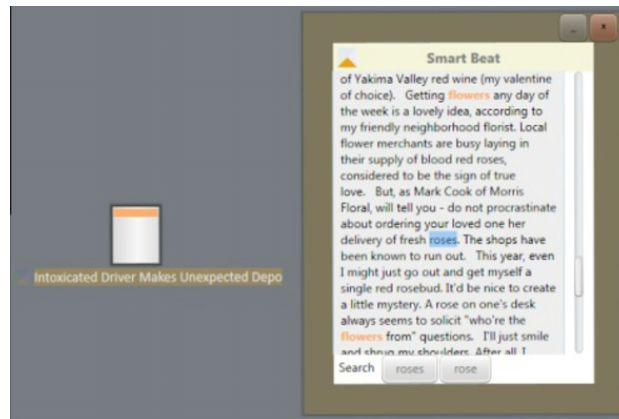
Fig. 11. In addition to the minimal document representation shown within a search result box, Cambiera also scales documents up to a "closed" (left) and "open" (right) state. This task-specific variation in sizes affords clutter reduction.
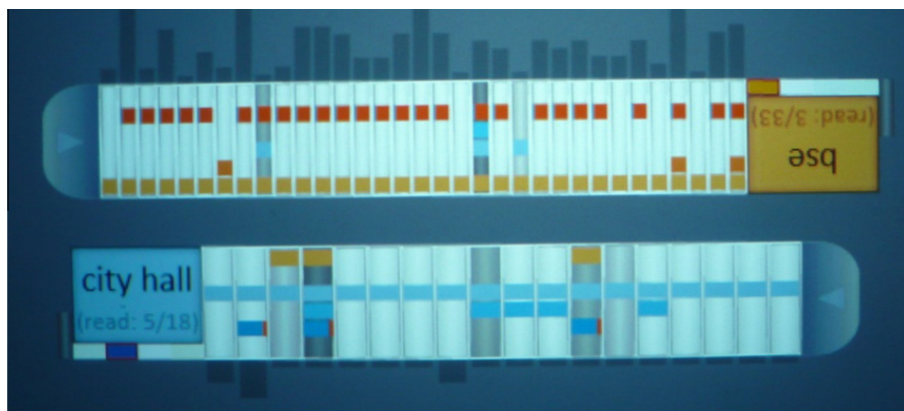
Fig. 12. Each partner using Cambiera is associated with a different color scheme (blue and orange, in this example). Collaborative brushing and linking provides awareness of how partners' findings overlap. For example, the fourth document under "city hall" has been read several times, and was found by "BSE", "city hall", and several other queries initiated by the blue user. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

position. However, a special gesture can be used to transfer document ownership to one's partner, which results in automatic re-orientation of that document.

3.4. WeSearch

Our final exemplar system, WeSearch (Morris et al., 2010), is designed to support groups of up to four users conducting collaborative Web search tasks, including the accompanying sensemaking process. For example, WeSearch might be used by a family to trying to decide what type of laptop computer to purchase for their home, or by a group of colleagues selecting a location for their company's next offsite meeting. WeSearch is designed for the custom-built, large-form-factor tabletop used by FourBySix Search (but with no augmentation by physical keyboards) (Fig. 13). Table 2 shows how WeSearch fits into our design space.

Each group member has a color-coded toolbar containing a search box. Touching the search box opens a virtual keyboard that can be used to enter queries. Search engine results for that query are then displayed in a custom-designed Web browser, oriented to face in the direction of that user's toolbar. In order to minimize clutter and facilitate sensemaking, the WeSearch browser offers the ability to divide a Web page into a set of *clips* by parsing the page's HTML (Fig. 14). Each clip is a small chunk of content (e.g., a single paragraph or a single image), that can then be removed from the browser; the originating Web page can then be closed to preserve space on the tabletop, leaving only the chosen clips behind.

In addition to a search box, each group member's toolbar also contains a *marquee* (Fig. 15). The marquee is designed to facilitate awareness, by showing color-coded content from other group members (such as query terms they have used and the titles of Web pages they have opened); this content slowly scrolls past in the marquee (the marquee can also be scrolled manually, in order to review or re-find content). To reduce tedious virtual-keyboard text entry, terms appearing in the marquee can be re-used by dragging them out of the marquee region and dropping them on the search box.

A button on the toolbars allows users to create *containers*, such as lists, grids, or free-form canvases. These containers can be used to collect and organize related clips, in order to support sensemaking (Fig. 16). The ability to add labels to clips
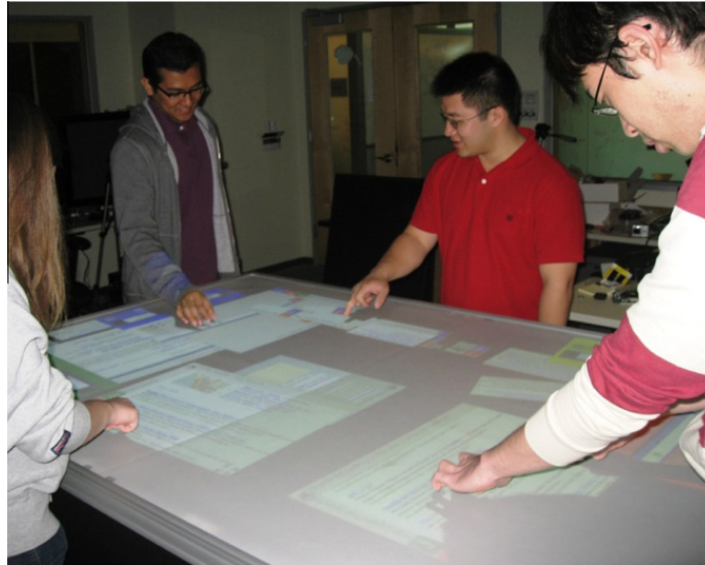
Fig. 13. A group uses WeSearch for a collaborative Web search and sensemaking session. WeSearch uses a custom, standing-height tabletop that is 4 feet wide and 6 feet long.
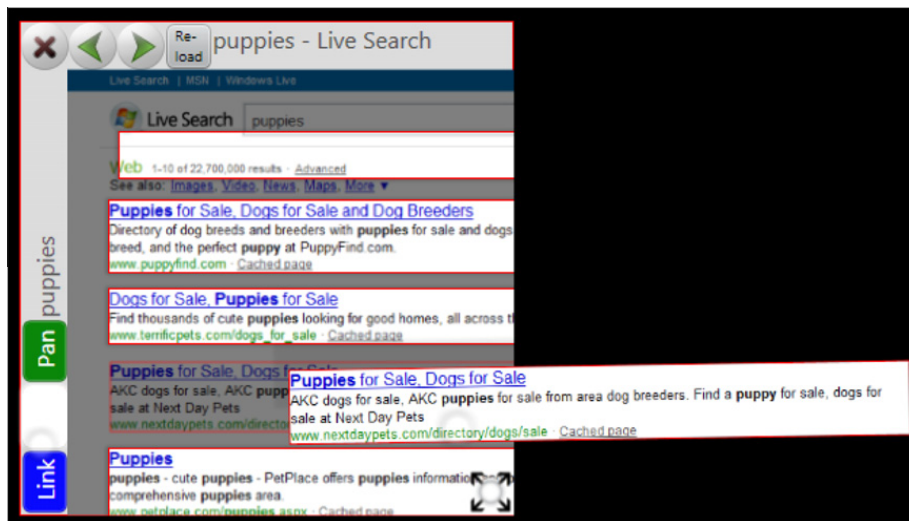

Fig. 14. The WeSearch Web browser offers the ability to automatically subdivide a Web page into clips based on the page's HTML, enabling removal of individual paragraphs or images by dragging the desired clip outside the boundaries of the browser.
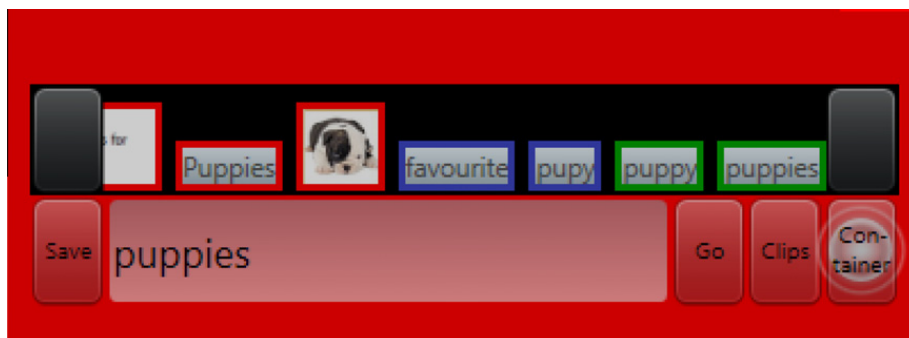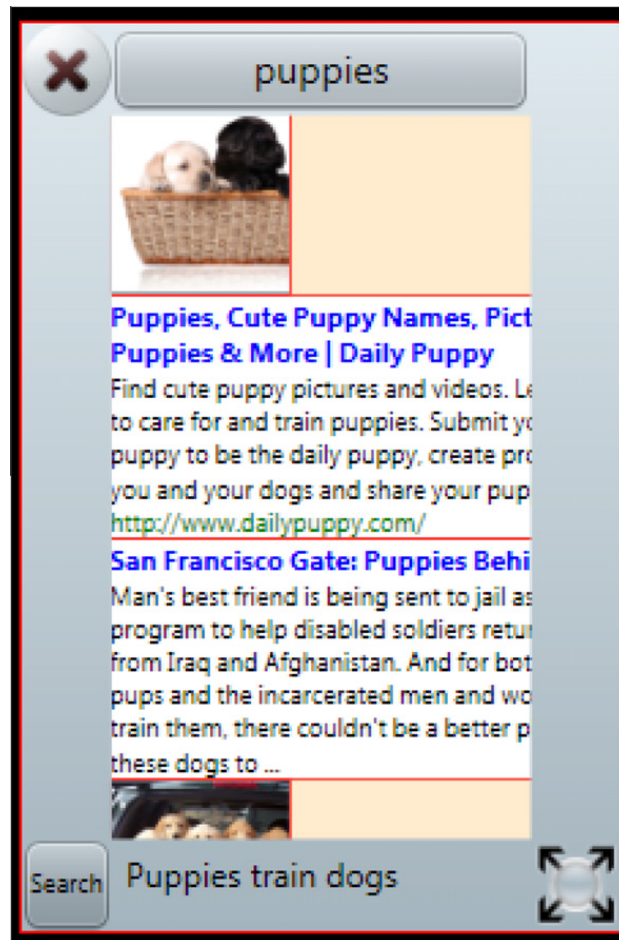

Fig. 15. Each group member has their own toolbar, consisting of a marquee (top half) and a search box (bottom half). Color coded items slowly scroll along the marquee, providing awareness of the activities of other group members, including query terms, page titles, and clips. These items can be dragged from the marquee to the search bar in order to facilitate text re-use.

Fig. 16. Users can create containers to organize their clips in various ways (in a list, in this example). Clips can be added to and removed from containers via drag-and-drop. WeSearch analyzes the contents of each container in order to suggest related query terms (previewed along the container's bottom border); pressing the container's search button causes this "search by example" query to be executed.

further supports the sensemaking process. WeSearch analyzes the set of clips stored in each container, and uses the space along each container's bottom border to suggest query terms that are thematically related to the clips using several heuristics, such as common terms shared by the set. This "search by example" feature (Fig. 16) can further reduce the need for soft-keyboard text entry. Users can also continue their sensemaking process away from the tabletop by using WeSearch's export feature, which creates an XML file (and accompanying XSLT file) that records information about the currently open clips (such as the clip contents, originating Web page, and the group member who created the clip), and information about their arrangement in containers; this file can then be viewed in any Web browser, such as on a user's own PC or mobile phone.

WeSearch approaches the challenge of text entry on tabletops from two angles: by facilitating re-use of other group members' query terms and terms from Web pages via the marquee, and by suggesting potentially useful query terms via the search-by-example feature. In addition to employing a large-form-factor tabletop, WeSearch attempts to reduce clutter by enabling users to extract clips from a Web page; in addition to conserving space, selecting pertinent clips can also facilitate the sensemaking process, by indicating what information is most relevant to the group. Finally, WeSearch addresses the orientation challenge by providing each group member with a personal toolbar containing the marquee, which provides awareness of the content being read by other group members at an appropriate viewing angle for each user.

4. Discussion

TeamSearch, FourBySix Search, Cambiera, and WeSearch paint four pictures of how interactive tabletops can support collaborative search tasks; each supports different task domains and each offers different ways of addressing the challenges of text entry, clutter, and orientation. Table 2 illustrates how our own and others' systems cover partially overlapping portions of the table + search design space – while each system covers some distinct aspects of the space (such as the target application domain, and the collaboration style supported), there is also a great deal of overlap, which means that some aspects of the design space are left unexplored. These unexplored aspects of the design space highlight important areas for future work.

For example, these systems focus mainly on supporting the information retrieval process, but sensemaking and report preparation are also key parts of many search tasks. Designing systems that take advantage of tabletops' spatial layout

and high awareness properties for sensemaking (or for integrating information retrieval, sensemaking, and reporting activities together), is a high-value area for further exploration.

These systems are also focused on a table-only device ecology (while FourBySix Search incorporates keyboards, we consider this a table-only system since the keyboards neither contain nor symbolize data). Tabletop researchers have explored the potential for synergy between tables and handheld devices (Wilson & Sarin, 2007), laptops (Shen, Everitt, & Ryall, 2003), and "smart" rooms (Johanson et al., 2002); however, the utility of this larger ecology of input and output devices for facilitating collaborative search on tables remains unexplored. Such additional devices could enrich the search experience, such as by facilitating text entry using a phone's keypad, allowing high-resolution reading on a secondary monitor, or enabling personalized search results via combining techniques such as groupization (Morris, Teevan, & Bush, 2008) with the data repositories on personal devices. Auxiliary devices can also be used for pre-seeding a tabletop search session with the outputs of earlier, private search sessions, or for capturing the state or results of a shared tabletop search for portability to other locations and work contexts.

Exploring the potential of tabletops to support different group configurations is also an exciting area for further work, such as supporting large groups (in classrooms, meetings, or public spaces). Tabletops, though often thought of as devices for co-located work, might also provide an interesting platform for mixed-presence or remote collaborations.

5. Conclusion

Interactive tabletops offer the potential to help migrate traditional collaborative practices, such as meeting face-to-face around tables and conducting joint research using traditional libraries and media, into the digital age, by serving as an engaging platform for collaborative search applications. In this article, we articulated the affordances that make tabletops a compelling platform for collaborative search, as well as the challenges they present in such circumstances. We described four prototype applications, TeamSearch, FourBySix Search, Cambiera, and WeSearch, that illustrate the potential of tabletops for collaborative search, and showed the different approaches each takes to addressing the challenges of supporting search on tables. By considering the larger design space of possibilities when combining collaborative search and interactive tables, we can see what areas have been explored by our own and others' forays into this space, and what areas remain unexplored. While our initial prototypes are promising, it is clear that the potential of this form-factor to support collaborative information seeking has only begun to be explored; we eagerly anticipate that many exciting advances will arise at the intersection of these two fields!

Acknowledgements

References

Amershi, S., & Morris, M. R. (2008). CoSearch: A system for co-located collaborative web search. *Proceedings of CHI*, 1647–1656.
Benko, H., Morris, M. R., Brush, A. J. B., & Wilson, A. (2009). Insights on interactive tabletops: A survey of researchers and developers. *Microsoft Research Technical Report* (March).
Blackwell, A., Stringer, M., Toye, E., & Rode, J. (2004). Tangible interface for collaborative information retrieval. *Proceedings of CHI*, 1473–1476.
Buja, A., McDonald, J. A., Michalak, J., & Stuetzle, W. (1991). Interactive data visualization using focusing and linking. *Proceedings of IEEE Visualization*, 156–163.
Dietz, P., & Leigh, D. (2001). Diamond touch: A multi-user touch technology. *Proceedings of UIST*, 219–226.
Fidel, R., Bruce, H., Peitersen, A., Dumais, S., Grudin, J., & Poltrock, S. (2000). Collaborative information retrieval. *Review of Information Behavior Research, 1*(1), 235–247.
Haller, M., Brandl, P., Leithinger, D., Leitner, J., Seifried, T., & Billinghurst, M. (2006). Shared design space: Sketching ideas using digital pens and a large augmented tabletop setup. *Proceedings of ICAT*, 948–959.
Han, J. (2005). Low-cost multi-touch sensing through frustrated total internal reflection. *Proceedings of UIST*, 115–118.
Hansen, P., & Jarvelin, K. (2005). Collaborative information retrieval in an information-intensive domain. *Information Processing and Management, 41*(5), 1101–1119.
Hartmann, B., Morris, M. R., Benko, H., & Wilson, A. (2009). Augmenting interactive tables with mice & keyboards. *Proceedings of UIST*, 149–152.
Hinrichs, U., Hancock, M., Collins, C., & Carpendale, S. (2007). Examination of text-entry methods for tabletop displays. *Proceedings of IEEE Tabletop*, 105–112.
Isenberg, P., & Fisher, D. (2009). Collaborative brushing and linking for co-located visual analytics of document collections. In *Proceedings of eurographics/IEEE-VGTC symposium on visualization (EuroViz 2009)* (pp. 1031–1038).
Izadi, S., Agarwal, A., Criminisi, A., Winn, J., Blake, A., & Fitzgibbon, A. (2007). C-slate: A multi-touch and object recognition system for remote collaboration using horizontal surfaces. *Proceedings of IEEE Tabletop*, 3–10.
Johanson, B., Fox, A., & Winograd, T. (2002). The interactive workspaces project: Experiences with ubiquitous computing rooms. *IEEE Pervasive Computing, 1*(2), 67–75 (April–June 2002).
Kruger, R., Carpendale, S., Scott, S. D., & Greenberg, S. (2003). How people use orientation on tables: Comprehension, coordination, and communication. *Proceedings of GROUP*, 369–378.
Large, A., Beheshti, J., & Rahman, T. (2002). Gender, differences in collaborative web searching behavior: An elementary school study. *Information Processing and Management, 38*(3), 427–443.
Maekawa, T., Hara, T., & Nishio, S. A. (2006). Collaborative web browsing system for multiple mobile users. *Proceedings of PERCOM*, 22–35.
Morris, M. R. (2008). A survey of collaborative web search practices. *Proceedings of CHI*, 1657–1660.

Morris, M. R., Brush, A. J. B., & Meyers, B. A. (2008). Field study of knowledge workers' use of interactive horizontal displays. *Proceedings of IEEE Tabletop*, 113–120.

Morris, M. R., Lombardo, J., & Wigdor, D. (2010). WeSearch: Supporting collaborative search and sensemaking on a tabletop display. *Proceedings of CSCW 2010.*

Morris, M. R., & Horvitz, E. (2007). SearchTogether: An interface for collaborative web search. *Proceedings of UIST*, 3–12.

Morris, M. R., Paepcke, A., & Winograd, T. (2006). TeamSearch: Comparing techniques for co-present collaborative search of digital media. *Proceedings of IEEE Tabletop*, 97–104.

Morris, M. R., Paepcke, A., Winograd, T., & Stamberger, J. (2006). TeamTag: Exploring centralized versus replicated controls for co-located tabletop groupware. *Proceedings of CHI*, 1273–1282.

Morris, M. R., Teevan, J., & Bush, S. (2008). Enhancing collaborative web search with personalization: Groupization, smart splitting, and group hit-highlighting. *Proceedings of CSCW*, 481–484.

Paek, T., Agrawala, A., Basu, S., Drucker, S., Kristjansson, T., Logan, R., et al. (2004). Toward universal mobile interaction for shared displays. *Proceedings of CSCW*, 266–269.

Paul, S., & Morris, M. R. (2009). CoSense: Enhancing sensemaking for collaborative web search. *Proceedings of CHI*, 1771–1780.

Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P., & Back, M. (2008). Algorithmic mediation for collaborative exploratory search. *Proceedings of SIGIR*, 315–322.

Rogers, Y., & Lindley, S. (2004). Collaborating around large interactive displays: Which way is best to meet? *Interacting with Computers*.

Russell, D. M., Stefik, M. J., Pirolli, P., & Card, S. K. (1993). The cost structure of sensemaking. *Proceedings of CHI*, 269–276.

Shen, C., Everitt, K., & Ryall, K. (2003). UbiTable: Impromptu face-to-face collaboration on horizontal interactive surfaces. *Proceedings of UbiComp*, 281–288.

Shen, C., Lesh, N., Vernier, F., Forlines, C., & Frost, J. (2002). Building and sharing digital group histories. *Proceedings of CSCW*, 324–333.

Shen, C., Vernier, F., Forlines, C., & Ringel, M. (2004). DiamondSpin: An extensible toolkit for around-the-table interaction. *Proceedings of CHI*, 167–174.

Smeaton, A. F., Lee, H., Foley, C., & McGivney, S. (2006). Collaborative video searching on a tabletop. *Multimedia Systems Journal, 12*(4), 375–391.

Stewart, J., Bederson, B., & Druin, A. (1999). Single display groupware: A model for co-present collaboration. *Proceedings of CHI*, 286–293.

Sun, X., Chiu, P., Huang, J., Back, M., & Polak, W. (2006). Implicit brushing and target snapping: Data exploration and sense-making on large displays. *Proceedings of AVI*, 258–261.

Talja, S. (2002). Information sharing in academic communities: Types and levels of collaboration in information seeking and use. *Review of Information Behavior Research*, 143–159.

Tang, A., Tory, M., Po, B., Neumann, P., & Carpendale, S. (2006). Collaborative coupling over tabletop displays. *Proceedings of CHI*, 1181–1190.

Tuddenham, P., & Robinson, P. (2009). Territorial coordination and workspace awareness in remote tabletop collaboration. *Proceedings of CHI*, 2139–2148.

Twidale, M., Nichols, D., & Paice, C. (1997). Browsing is a collaborative process. *Information Processing and Management, 33*(6), 761–783.

von Ahn, L., & Dabbish, L. (2004). Labeling images with a computer game. *Proceedings of CHI*, 319–326.

Wigdor, D., & Balakrishnan, R. (2005). Empirical investigation into the effect of orientation on text readability in tabletop displays. *Proceedings of ECSCW*, 205–224.

Wilson, A., & Sarin, R. (2007). BlueTable: Connecting wireless mobile devices on interactive surfaces using vision-based handshaking. *Proceedings of Graphics Interface*, 119–125.