

Designing Conversation-context Recommendation Display to Support Opportunistic Search in Meetings

Nan Li
EPFL
Switzerland
nan.li@epfl.ch

Pierre Dillenbourg
EPFL
Switzerland
pierre.dillenbourg@epfl.ch

ABSTRACT

In meetings, people sometimes come across information in the environment or in the conversation. This kind of accidentally encountered information may be or not be relevant to the main meeting topic, but stirs people's temporary interest and mediates group conversation. In some cases people even need to search these unexpected terms in the Web. This phenomenon is called opportunistic search. We question if digital displays can be designed to support opportunistic search in meetings. Assuming information sharing in collocated meetings may largely come from conversation, we propose to design an ambient tabletop display that provides just-in-time conversational information to support opportunistic browsing and searching. In this paper, we present our design explorations with user study, and answer some of the major design questions for such systems.

Categories and Subject Descriptors

H.5.2 [Information interfaces and presentation (e.g., HCI)]: Miscellaneous.

General Terms

Human Factors; Design.

Keywords

Ambient display; context-awareness; speech recognition; information retrieval; interactive tabletop

1. INTRODUCTION

A meeting setting is usually temporarily created by people gathering around for a certain task, forming a social atmosphere that can stir serendipitous sharing of information. During the process of information sharing, the need of querying the Web often arises. Depending on meeting types, search frequency and intensity may vary to different extent. In information-seeking meetings (e.g. shopping

for new products, trip planning), searching is a goal-driven and purposeful activity, so that collaborative search tools must be available and people search both intentionally and intensively. In other meeting scenarios where seeking information is not the main goal (e.g. strategy decision making, brainstorming), searching is usually a temporary sub-activity which occurs opportunistically, but last for short periods. This is what we refer by the term "opportunistic search" in this paper. For example, people may spontaneously seek information for certain terms that they encounter unexpectedly either from meeting conversation or resources (paper documents or other objects in the surroundings) available in the meeting environment. Such Web search activities mainly serve for two purposes, problem clarification and idea inspiration. The former helps a group reach mutual understanding about a topic, and in the latter case Web search may inspire new ideas when people are engaged in a discussion that requires creativity, e.g. brainstorming.

Some aspects of opportunistic activities in meetings have been studied in UbiComp and CSCW. De Bruijn et. al presented a coffee table with an information display [1] where people may discover useful information serendipitously. Zancanaro et.al showed how a table that displays museum related information may affect people's discussion topics while drinking coffee [5]. Both systems are supposed to serve casual meetings without specific tasks, and the displays are not interactive, but just act as information portal. The discovery of information for people sitting around the former coffee table is totally opportunistic, since the information shown is randomly chosen. For the latter case, the information presented is adapted to social dynamics such as participants' attention, but the information contents are constrained with museum topics, which are not adapted to group's conversation context. WordPlay [2] proposed a tabletop system that can semantically associate conversational words to help brainstorming. However, the system is not considered as ambient display but rather a brainstorming supporting tool. These research projects concern more about the effects of opportunistic information browsing instead of information seeking in meeting spaces. Moreover, to our knowledge, none of the aforementioned systems have been formally evaluated.

In meetings where opportunistic searching occurs, collaborative search facilities [3] are not necessarily required due to the system's high complexity. In current meeting practices, people can access information with personal digital devices, e.g. smart phones, tablets or laptops. This requires temporary attention shift from the meeting task to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MUM '12, December 04 - 06 2012, Ulm, Germany

Copyright 2012 ACM 978-1-4503-1815-0/12/12 ...\$15.00.

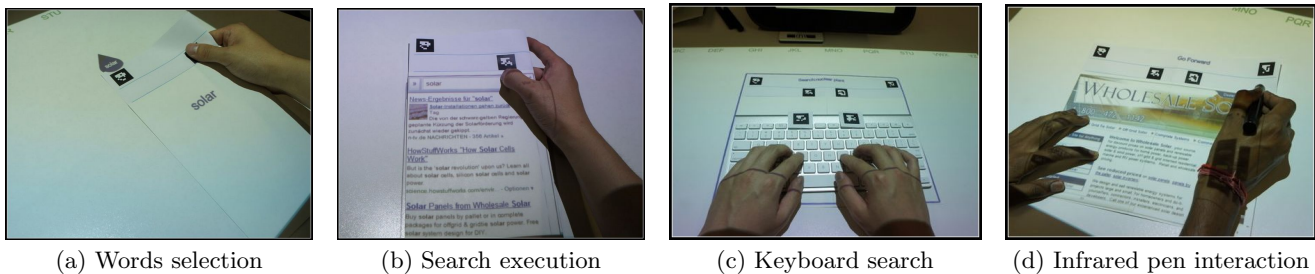


Figure 1: Paper Browser Interactions

a search task. Sometimes people may even be reluctant to make a Web search, due to impoliteness, shyness and fear to show ignorance, or when their subjective perception of the effort/reward ratio for making that search is low. The AMIDA ACLD [4] proposed a solution of automatically searching and linking documents within conversational context, which made free of explicit search tools. However, this system also lacks user evaluations, and we think the text-based swing interface is not suitable for meeting settings. Given optimal speech recognition technology, we question how to design user interface to support opportunistic search activities during a meeting by visualizing conversation context in an unobtrusive way. The system should ease user experience of accessing and retrieving information. In this paper, we will present our explorations for the design of opportunistic search meeting support systems with tabletop display.

2. QUERY-TERM SUGGESTION INTERFACE

One of the major design choices we faced when designing the initial prototype was that we did not know whether or not to provide users with Web search results that are automatically retrieved based on their conversation context or to give them the freedom to search with selected conversational words. As little literature has explored opportunistic search in meetings, we decided to start with the second option, which seemed to be a more intuitive way of querying the Web. We call this design "query-term suggestion interface" and we want to answer three questions through user experiments in order to guide future designs.

- Is building search queries from conversational words effective and efficient for opportunistic search?
- Does the system really work unobtrusively?
- Is there a need to give each user a separate display in order support parallel search and browsing?

2.1 System Design

We use a portable projection-camera system, the Tinker-Lamp [6] to yield an interactive surface for the ambient display. We adopt raindrop metaphor for the design of the display. Words are enclosed in rain drops. These drops containing just-in-time conversational words are vertically ordered by their arrival time, and each drop moves down slowly from the upper border and disappears at the bottom. A rain drop gets highlighted when the contained word is repeatedly spoken. The words are ordered horizontally according to their first letter's alphabetic order. As meeting goes

on, the display is filled with word drops in an arranged manner (See Figure 2.), in order for people to easily spot them on the display. The tabletop surface continuously displays




Figure 2: The Query-term Suggestion Interface

conversational words. When a user comes across something that she is willing to search, she can use a foldable paper browser to intercept a dropping word with her thumb pressing on one marker and holding the paper below the drop (Figure 1(a)). Multiple words can be selected with the same gesture. By releasing the thumb, results are shown in the browser (Figure 1(b)). The system also supports traditional search style by typing words with a keyboard. Each meeting participant has a paper browser, but only one keyboard is available for the group. Simply by placing the keyboard onto a paper browser, a user can even type her own query terms for search (Figure 1(c)) and then use an infrared pen to interact with the browser content (Figure 4(d)). Parallel search with keyboard is not supported, but different users can browse at the same time. Different paper browsers share the same link-visit-history, which is listed on the back side of each paper sheet, but displayed contents are not synchronized. We use paper as additional personal displays because they can be easily "plugged" into and out of the tabletop display, which separates the ambient display and the interaction space without altering the existing environment.

2.2 User Study

Several experiments were conducted to obtain qualitative usability feedback about the system. We recruited 12 users (4 groups, each containing 3 users) to participate our study. We gave each team a map and a set of power plant technologies and asked them to reach an agreement in 30 minutes on which technologies to use for building power plants for a given city. Considering the performance of real speech recognition technology might bias the evaluation, we decided to use Wizard-of-OZ method by employing a person to over-

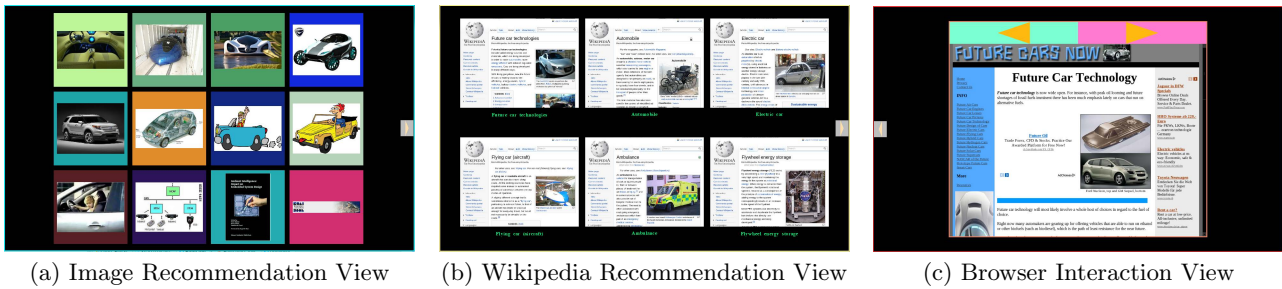


Figure 3: The Search Result Suggestion Interface

hear the meeting conversation and enter spoken words into the system. The word-filtering rule is that original forms of nouns, verbs, adjectives and adverbs occurred in the conversation, semi-structured interviews, and post-test questionnaires.

According to our in-situ observations of the experiment, we found that the need for information seeking occurred several times in every group meeting. Most of the Web searches were made with keyboard, although only one keyboard was available for each group. Users seldom used paper browsers separately simultaneously. At the beginning of each meeting, the users looked at the display quite often, perhaps due to novelty effect of the display. During this period, they had tried to pick up some words on the table. Mostly users were purposefully looking for words that they have just spoken in the past few seconds. If the required words could not be found, then they usually spoke them out again loudly and then waited for the words to appear (e.g. “*tide power station, where is tide? tide! tide!*”). However, such behaviors did not last for long. Most of the groups forgot the word suggestion feature after playing with it for a while. They started to search in a traditional way, i.e. directly with a keyboard. In the questionnaire, all users reported search activities during the meeting, because they needed more information to complete the task. We also asked if the display had provided good key words for search. 7 out of 12 users gave negative answers. In fact, the users all knew very well how to query a search engine. Keyboard was much more intuitive than the word selection interaction with the paper browser in this situation, especially when they knew clearly what to search for. No one reported having searched for something due to inspiration of certain words shown in the display.

2.3 Discussion

Back to the questions we posed at the beginning of this section, the experiment has revealed that the idea of displaying conversational words for forming search queries is problematic, but we have learned some lessons. It is difficult for the users to construct search queries by picking up words from a display which is full of words. This would increase the cognitive load because users have to both forming the query and looking for words on the table at the same time. When people already have something in mind, they would rather use keyboard to type the query terms. Simply displaying words failed in increasing the chance of opportunistic search, since a single word conveys little information so that it is often ignored by people engaging in a seri-

ous cognitive-demanding task. Animation of the raindrops were reported disturbing in early periods of the meeting, but as users had lost trust and interest, the displayed information totally goes to the background. This indicates that animations are not necessary in attracting users but the displayed information content does. Although sometimes they had parallel browsing behavior, multiple-browsers support is definite not a design requirement. “When to search what and by who” can often be regulated by social factors in the meeting.

3. SEARCH RESULT SUGGESTION INTERFACE

We soon developed a completely new design, where we abandoned multiple displays, unnecessary animations, and even the idea of suggesting query-terms from conversation. We decide to take the other design choice mentioned at the beginning of last section, i.e. handing over the Web querying process to the machine, and the ambient display presents up-to-date search results only. We posed two questions to be answered by a user study.

- Are media contents (images) more attractive than textual contents (Wikipedia) for servicing ambient information?
- Do users still need to search with their own keywords with the presence of automatic search results?

3.1 System Design

We use TinkerLamp and a Microsoft Kinect sensor which can detect finger touches on the table to build an interactive ambient display. The system takes every N consecutive conversational keywords to query Microsoft Bing and show the top M image results or Wikipedia results. Each search result stays for at least 5 seconds. When new results arrive, the display is updated. For parameters we choose $N = M = 3$, since we find that with these values, the refreshing rate of the display can keep pace with the conversation in our Wizard-of-OZ configurations. The ambient display has two views, a recommendation view (Figure 3(a), 3(b)) and a browser interaction view (Figure 3(c)). The recommendation view displays up-to-date search results in tabular layout to increase the possibility of opportunistic discovery. If a user is interested in one of the displayed information, she can select it with hand. Then the display will automatically switch to the browser interaction view, which contains a fully-functional multi-touchable Web browser that automatically redirects to the selected link. There is also a rectangular arrow button

on each view to allow manually switching between different views.

3.2 User Study

We again recruited twelve subjects (4 groups, each containing 3 users) to participate our study. The experiment had a within-group design, with each group completing two open-ended brainstorming tasks using two different search suggestions, with one task using Wikipedia and one using images. The order of tasks and suggestion types was balanced using a Latin Square design. We used two brainstorming tasks: house of the future and car of the future. For each task, 7 minutes was given for a divergent thinking phase to develop ideas in the team without judgment, and this is followed by a 8 minutes convergent thinking phase when they must shortlist the ideas according to their feasibility within twenty years. We collected data by in-situ observations, semi-structured interviews, user interaction logs, and questionnaires.

For each task that the users have done, the display suggested 108 - 221 results in total, depending on the amount of speech. From the interaction log we found that the average usage rate of the display is 3.75 times per experiment. The usage rate only reflects how many times they have interacted with the display (read details in the browser interaction view), not how many times they had glanced the display. From our in-situ observation, the ambient display attracted more attention from participants than the previous design, but received less complaints about disturbance during the whole experiment. The users discussed about the suggested results many times, including Wikipedia pages they were looking for and unexpectedly interesting images. Actually, since the users were having a brainstorming task, they might be opportunistically inspired by the images or Wikipedia pages shown in the recommendation view, without the need of knowing more details in the browser interaction view. Some users even intentionally looked for information on the display when they ran out of ideas (e.g. "...*what else ideas....ah!...let's see what the display suggests...*"), while many others clicked the search results due to curiosity and interests (e.g. "*I came cross an interesting image, so I just wanted to see more information..but..it actually has nothing to do with the task*"). When participants were asked to compare the two designs subjectively, people tent not to agree that image caught more of their attention and were more useful than Wikipedia results both in divergent phase (3.5 in 7-point Likert scale) and convergent phase (3.75 in 7-point Likert scale). The results are almost neutral, since users did not have the impression that either one clearly outperformed the other, and both views have attracted them equally. An interesting finding from collected subjective feedback is that people had less tendency to search for other information with customized keywords in the image condition than in the Wikipedia conditions (3.25 ~ 4.92, $F(1, 11) = 7.86, p = .017$).

3.3 Discussion

When analyzing the results, we do not much focus on which design has led the group to create more ideas in the discussion, since in our experiment it is difficult to measure objectively how their finalized idea list may correspond to what had been suggested on the display. We are then more interested in how the different designs of the display me-

diates group interaction. If we revisit the two questions posed before, the experiment has shown that users have no preference towards either design. It is not possible for a machine to automatically and correctly get users' every search need, so users still may want to search by their own. Both interface designs have serviced information in their own rights, but textual informational interface seemed significantly more persuasive in inspiring people to search with their own keywords.

4. CONCLUSION

In this note we presented two evolving tabletop designs for opportunistic search support in meetings with corresponding user studies. We aimed at exploring design possibilities rather than proposing design guidelines. The research itself has some limitations, e.g. we had limited number of users for experiments, and used different tasks for two systems. Nonetheless, it still provided implications for designing future systems. The failure of the query-term interface does not reflect the negation of suggesting keywords, it actually suggests against "animated words" due to their visual disturbance. We find out that the need of searching on customized words is inevitable, so simple search user interface should be provided. For example, intelligently extracted keywords can be displayed in a weighted word clusters, allowing users to build customized queries out of them. The display should also present rich media representation of the conversation. Full Wikipedia pages do not outperform their image counterparts in terms of recommendation media, so displaying fullpage is a waste of display resource. We plan to employ a real speech recognition system in our next step to see its potential and constraints. Contextual information other than speech is also considered to be used in generating recommendations, such as hand-writings, drawings and digital text-input. Mobile phones can be used as private information space to interact with the recommendation display. We envision such ambient displays will become ubiquitous in future meeting spaces.

5. REFERENCES

- [1] O. d. Bruijn and R. Spence. Serendipity within a ubiquitous computing environment: A case for opportunistic browsing. In *Proc. UbiComp*, pages 362–370, 2001.
- [2] S. Hunter and P. Maes. Wordplay: A table-top interface for collaborative brainstorming and decision making. *Proc. IEEE*, 2008.
- [3] M. Morris, J. Lombardo, and D. Wigdor. Wesearch: supporting collaborative search and sensemaking on a tabletop display. In *Proc. CSCW*, pages 401–410, 2010.
- [4] A. Popescu-Belis, E. Boertjes, J. Kilgour, P. Poller, S. Castronovo, T. Wilson, A. Jaimes, and J. Carletta. The amida automatic content linking device: Just-in-time document retrieval in meetings. In *Proc. MLMI*, pages 272–283. 2008.
- [5] M. Zancanaro, S. Oliviero, D. Tomasini, and F. Pianesi. A socially aware persuasive system for supporting conversations at the museum. In *Proc. IUI*, pages 395–398, 2011.
- [6] G. Zufferey, P. Jermann, and P. Dillenbourg. A tabletop learning environment for logistics assistants: activating teachers. In *Proc. IASTED HCI*, pages 37–42, 2008.