Citation for published version:

Kelly, R & Payne, SJ 2014, 'Collaborative web search in context: a study of tool use in everyday tasks', Paper presented at 2014 ACM Conference on Computer Supported Cooperative Work, Baltimore, USA United States, 15/02/14 - 19/02/14 pp. 807-819. <https://doi.org/10.1145/2531602.2531617>

DOI:

[10.1145/2531602.2531617](https://doi.org/10.1145/2531602.2531617)

Publication date:

2014

Document Version

Peer reviewed version

[Link to publication](#)

© ACM 2014. This is the author's version of the work. It is posted here for your personal use. Not for redistribution

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Collaborative Web Search in Context: A Study of Tool Use in Everyday Tasks

Ryan Kelly

Department of Computer Science
University of Bath, United Kingdom
r.m.kelly@bath.ac.uk

Stephen J. Payne

Department of Computer Science
University of Bath, United Kingdom
s.j.payne@bath.ac.uk

ABSTRACT

Recent research efforts have led to the creation of a number of systems that provide specialised support for collaborative web search. However, the use of these tools has not been studied outside of the laboratory, and as collaborative search becomes increasingly commonplace in everyday life, there is a need to understand whether the support provided by collaborative search systems fits with real-world information seeking practices. In the present study, we deployed two collaborative search tools to pairs of searchers with genuine information needs. We report findings from in-depth interviews conducted after searchers had used their assigned system for an extended period of time. Our findings show how system features were used and appropriated in pursuit of collaboration, throwing light on the way in which collaborative search is conducted in quotidian settings. Theoretical and practical implications of the results are discussed.

Author Keywords

Collaborative information seeking; Collaborative search

ACM Classification Keywords

H.5.3 Information Interfaces and Presentation (e.g. HCI):
Collaborative Computing

INTRODUCTION

Information seeking on the web often occurs as a collaborative act between two or more searchers with a shared information need. Research has revealed that such behaviour is fairly prevalent [3, 14] and appears to be on the increase [15]. In response to these findings, a number of prototype systems offering specialised support for collaborative search have been developed within the research community (e.g. [6, 12, 13, 16, 17, 18]) and some commercial systems (e.g. *SearchTeam*, *Bing*) now provide support for collaborative and other socially-oriented search behaviours [15].

Despite this progress, the actual success of collaborative search systems, in terms of mainstream takeup, has been fairly limited. To the best of our knowledge, none of the tools presented in the

research literature has achieved widespread adoption. Morris [15] further notes that the majority of commercial systems are either defunct (e.g. *Aardvark*, *Flock*) or remain in the early stages of development (e.g. *Pinterest*, *So.cl*). One potential reason for this is that existing collaborative search tools may require too much effort or do not offer meaningful benefit over ad hoc practices (e.g. link sharing via email) used during everyday tasks like travel planning and online shopping [14]. However, neither of these possibilities has been investigated. A related issue is that collaborative search tools have only been studied in relatively short-term laboratory evaluations. While such studies are fine for testing specific hypotheses and evaluating initial designs, they do not provide information about how systems are used over the longer term. No work has, as yet, studied the efficacy of collaborative search tools in quotidian settings. It is, therefore, unclear as to how well existing tools fit with everyday collaborative search practices.

In the present study, we deployed two collaborative web search tools to pairs of searchers conducting everyday information seeking tasks. Our immediate aim was to understand whether the tools were useful in supporting actual collaborative search, and thus we recruited pre-established collaborators—with information needs that were intrinsically collaborative—tasks were self-selected and participants used their assigned system for as long as they wished, creating high external validity. By interviewing our participants about their experiences, our study allows us to understand how the tools were used in accordance with existing routines, in turn contributing to a broader theoretical understanding of collaborative search behaviour ‘in the wild’. And, by detailing how our participants used and appropriated particular system features, this research provides implications for the design of future tools to support everyday collaborative search tasks.

BACKGROUND & RELATED WORK

Within the broader area of social search, collaborative web search describes situations “in which participants work together to satisfy an information need” [15, pg.1182]. Such activity can occur either synchronously or asynchronously, and participants may be co-located or geographically distributed [9]. Collaborative search is usually scoped to consider explicit, intentional collaboration (cf. [8]), precluding consideration of filtering or recommendation tools that utilise prior searches from anonymous ‘collaborators’.

Several studies have revealed that collaborative search occurs during a variety of professional and personal tasks, including medical research, travel planning, and online shopping

[14, 15]. However, searchers report that managing such collaboration can be arduous, requiring workarounds such as link sharing via email [14, 15] or the use of ‘tools-at-hand’ (e.g., blogs, text documents) to keep track of search results [3]. Researchers have suggested that these behaviours can be supported at the user interface, and, to this end, a number of systems have been designed by the research community. Examples include *CoSearch* [1] *Results Space* [2], *Cerchiamo* [6], *Querium* [7], *Coagmento* [12], *ViGOR* [13], *SearchTogether* [16], *WeSearch* [17], and *CoSense* [18]. While the specific functionality offered by each tool is different, and can vary according to the anticipated scenario of use, the common goal has been to alleviate the need for workarounds by providing browser-based support for collaborative search. As an in-depth review of these systems is beyond the scope of this paper, it is instructive to focus on the general aspects of collaboration that such tools aim to support. These are *awareness*, *division of labour*, *persistence*, and *sensemaking* [5, 13, 16, 18].

Awareness refers to the ability to acquire knowledge about the current and past activities of one’s interaction partners [21]. Such knowledge promotes coordination and lessens the need for explicit communication about task progress [21]. Features for supporting awareness during collaborative search include shared query, browsing, and page visitation histories [2, 12, 16]; commenting of pages [12, 16]; and increased salience of particular results based on collaborator ratings [2].

Division of Labour refers to the process of distributing a task across members of a group [5]. The aim here is to facilitate concurrent work while preventing redundancy and duplication of effort. In collaborative search tools, division of labour has been supported through text chat systems, which allow collaborators to establish division of labour through communication [12]; automated splitting of search results [16]; allocation of tasks by role [6]; and algorithms for selective filtering [5].

Persistence, referring to the storage and display of activity from prior search sessions, enables asynchronous collaboration through information re-finding and resumption of prior search sessions [2, 16]. Persistence has been supported through retention of chat logs, pageview statistics, and session histories [12, 13, 16]; automatic generation of session summaries [16]; and relevance rating tools [2].

Finally, *sensemaking* support allows collaborators to understand the search process, in terms of *what* has been found, *how* it was found, and where tasks have been handed off between collaborators [18]. Example features include: context awareness through visualisation of search strategies and trajectories [18]; functions to exchange sections of webpages [17]; and browsable timelines of pages viewed by collaborators [12, 18].

While studies have suggested that the features described above are beneficial for supporting collaborative search, no work has yet examined the success of any tool when used in natural field settings. Instead, most have been studied using artificial tasks completed under the constraints of short-term laboratory evaluations. This is problematic because real-world search may involve protracted behavioural patterns that are not well supported by current systems. Additionally, a growing body of

empirical work suggests that various factors, including awareness [21], communication channel [11], time [10], and spatial proximity [20], can impact the collaborative search process. While such studies serve to advance theoretical accounts of collaborative search, they do not offer findings regarding the use of tools outside the lab.

Given the considerable research effort invested in supporting collaborative search, we believe that studying how current tools fare in everyday settings would be beneficial for directing future design activities. To this end, we designed a field study of collaborative search tool use, with the aims of gaining a better understanding of collaborative search in the wild and of identifying potential enhancements for future systems. We used two existing systems to study collaborative search behaviour. The fact that we did not develop either system means that we cannot access system logs containing quantitative data about tool use. This paper focuses on qualitative analysis of interviews conducted after our participants had used an assigned system to complete a real-world collaborative search task.

FIELD STUDY

System Choices

Our first system was *Coagmento*¹, a freely available tool that incorporates a range of features designed to support collaborative search (e.g. shared bookmarks, chat functionality) [12]. This tool can be regarded as a ‘general purpose’ collaborative search tool in that it is not intended to support any particular task over another. Since we did not specify in advance that our participants should engage in any particular type of search, this made *Coagmento* an appropriate choice for our study.

Our initial intention was to have all of our participants use *Coagmento*. However, a software update introduced a number of bugs (e.g. SQL database errors, malfunctioning UI elements) halfway through the study. These problems were beyond our control as we are not responsible for the creation of *Coagmento*. To avoid the problems affecting our remaining participants, we decided to switch to *Diigo*². While this choice was partly pragmatic (very few of the systems presented in the literature are available for outside use) *Diigo*’s overall functionality is actually very similar to *Coagmento*, meaning that a change of system did not require a large departure from our established methodology. And, as will be seen in our results, using a second system was beneficial in that differences between the two tools allowed us to obtain some valuable comparative insights.

System Functionalities

Both *Coagmento* (see Figure 1) and *Diigo* (Figure 2) are web-based systems. Each is comprised of two parts: a web browser plugin that provides rapid access to features intended for collecting, sharing, and saving information; and an online space to which users can save pages and view the results of prior search sessions. Both tools are also available as mobile applications, but these will not be discussed here as none of our participants opted to use them during our study.

¹<http://www.coagmento.org>

²<http://www.diigo.com>

Figure 1. Coagmento for Mozilla Firefox: (A) The Coagmento toolbar. (B) The CSpace with drop down selections for filtering saved results. (C) Shared history displaying thumbnails of recorded pages. (D) A larger preview of a page, accessed by clicking on the relevant thumbnail. (E) Information about a page's time of capture, project, and username of the original viewer. (F) Sidebar providing chat functionality, history, notifications, and a notepad. Image captured April 2013. Coagmento, Copyright © Chirag Shah. Image used with permission.

Figure 2. Diigo for Mozilla Firefox: (A) The Diigo toolbar. (B) Shared history with links to recorded pages. (C) Link to the user's personal library ('My Library'). (D) Sidebar providing a list of pages from the user's personal (but not group) library and a list of annotations on the current webpage. Image captured April 2013. Diigo, Copyright © Diigo Inc. Image used with permission.

Figure 1 shows the Coagmento system installed on a user's browser. The browser is open at the online 'CSpace' repository (B) to which pages are saved during search sessions. When the user is logged in to the system, all webpages visited are recorded to this history, appearing as thumbnails in chronological order (C). Each thumbnail can be clicked to reveal a larger preview of the relevant hyperlink (D), alongside the date and time of viewing and the username of the visitor (E). This history is then shared with all collaborators in a dedicated project folder, accessible via the CSpace.

The Coagmento plugin provides two components: a toolbar (Fig. 1, A) and a sidebar (F). The former includes buttons that allow users to bookmark whole pages, collect annotations and snippets, and recommend pages to collaborators. Pages captured using these tools are saved to the shared history and appear as an image bearing the relevant icon (the second thumbnail in area C shows the icon for a bookmarked page). Users can then selectively filter their history according to these differing types of content using the 'All Objects' dropdown (B). Of the remaining buttons, 'Resources' opens and closes the Coagmento sidebar, which contains a shared chat, a notepad, a history of recent bookmarks, and a notification submenu. 'Editor' provides access to a shared document for collaborative editing. Finally, the toolbar displays a summary of the current page in terms of views, snippets, and annotations [12].

Diigo (Fig. 2) provides a toolbar plugin (A) with functionality similar to Coagmento. Users are able to bookmark pages, leave highlights on relevant sections, or capture (screenshot) particular sections of a page. Each of these can then be saved to the Diigo webspace, either to a private library (accessed via the 'My Library' tab in area C of Fig. 2) or, in the case of bookmarks and highlights, to a shared 'group' space listing captured links in chronological order (Fig. 2, B). Of the remaining buttons, 'Send' allows users to email a page directly to a collaborator, and 'Read Later' allows for the webpage to be saved to a private library for later reading. The Diigo sidebar (D) provides quick access to the user's private library, a list of annotations for the user's current page, and a view of all prior readers of a particular webpage.

Comparison of Figures 1 and 2 reveals some differences which are worth considering due to their reference later in the paper. The first concerns the way in which each system saves and presents pages to users. In Coagmento, pages are represented using small thumbnails, each of which must be clicked to reveal further information about the represented page. In Diigo, no thumbnail is visible; instead, the title of the page is shown alongside a hyperlink and the username of the person who made the bookmark. Thus, the provision of initial descriptive information is different in each system—Coagmento relies solely on visual information, whereas Diigo uses text.

A second difference concerns the way in which each system tracks user behaviour. While active, Coagmento records *all* of the pages a user visits while logged in to the system, with individual pages represented by individual screenshots in the CSpace area (Fig. 1, C). This means that every page a user visits is captured by the system, regardless of whether or not the content is relevant to the user's primary information need.

In contrast, the only pages captured by Diigo are those that the user explicitly tells the system to save, either by bookmark, highlight, or capture (area B of Fig. 2 displays some examples). As will be seen in our results, this 'all or nothing' dichotomy leads to a number of concerns related to sensemaking and privacy during the display of shared search histories.

At this point, we should stress that it is not the aim of this paper to 'evaluate' Coagmento and Diigo in terms of their usability or relative successes and failures. Rather, our use of these systems is guided by the fact that they can be regarded as exemplary tools that might be used to support collaborative search—indeed, Coagmento has been designed in accordance with the research literature on collaborative search [12]. We have no investment in either system and no desire to demonstrate that one is better than the other. Our aim is to use these tools as probes to learn more about real-world collaborative search while identifying broader lessons for future systems.

METHOD

Study Design

Beyond our use of Coagmento and Diigo, we designed the present study to be as naturalistic as possible. Our first concern here was choice of tasks—this is a difficult yet critical issue when studying information seeking. As we were aiming for high external validity, we allowed participants to choose their own search tasks. This encourages intrinsic motivation while providing insights into real everyday search behaviour.

Second, the present study was open-ended and we did not impose any time constraints on our participants' search process. Not only did this provide ample time for our participants to explore and become familiar with system features, it allowed us to understand tool use during the broader collaborative search process, spanning tentative exploration through to results collection, refinement, and eventual selection.

Lastly, we allowed participants to use the systems according to their own preferences: at home, at work, or even on the move. We felt that learning about where, when, and how collaborative search occurs would help to understand how tools are used in the wild, and might also prove useful in terms of designing future technologies.

Participants

A total of 16 participants (eight pairs) took part in our study. Participants' ages ranged from 19–34 ($M = 24.4$, $SD = 4.7$). Pair 3 were both male; pairs 6 and 7 both female; and all other pairs were male/female. Pairs were comprised of friends or romantic partners, i.e. there were no anonymous pairings of unfamiliar participants. Participants were recruited via Facebook and our University noticeboard. Our adverts stated that we were looking for groups of people who would soon be completing a collaborative search task, offering the "opportunity to use a system designed to support collaborative information seeking behaviour". We used purposive sampling, vetting those who responded to our advert to ensure that their information needs were genuine. We did not have cause to turn away any of those who responded to our adverts. We offered each participant £20 as a goodwill gesture for completing the study.

Pair	System	Chosen task(s)	Distinct search sessions over time, by task	Pages saved, by task	Total duration of use
1	Coagmento	Travel planning	10 sessions over 14 days	27	14 days
2	Coagmento	Travel planning; House hunting	6 sessions over 16 days; 3 sessions over 8 days	10; 8	35 days
3	Coagmento	Concert venues	3 sessions over 12 days	4	12 days
4	Coagmento	Travel planning	3 sessions over 7 days	9	7 days
5	Diigo	Houses; Shopping	3 sessions over 5 days; 8 sessions over 3 days	3; 10	10 days
6	Diigo	Houses; Travel planning; Shopping	15 sessions over 14 days; 4 over 7; 3 over 3	18; 8; 14	21 days
7	Diigo	Houses; Travel planning	3 sessions over 7 days; 2 sessions over 8 days	22; 8	21 days
8	Diigo	Travel planning; Shopping	8 sessions over 6 days; 5 sessions over 8 days	10; 8	14 days

Table 1. Participants’ assigned system, task choices, distinct search sessions, number of items saved, and total duration of use. The ‘Pages saved, by task’ column refers to the total number of items captured using bookmarks, snippets, or annotations. Note that some tasks were completed concurrently and others were distinctly separate, meaning that aggregations of task completion do not equate to total duration of participation.

All participants were made aware from the outset that their payment was fixed and that it was not related to performance or time spent using their assigned system.

Table 1 provides information about the system assigned to each pair alongside task choices. Each pair used only one system and none changed tool during the study. The tasks chosen by our participants fit well with what the literature identifies about collaborative search [14, 15], and several pairs actually chose to complete more than one task during the study. While the nature of the information required by each of the tasks is different, we consider them to be qualitatively similar in that each is open-ended and allows for considerable latitude in terms of exploratory search behaviour. Furthermore, each involves an evolving information need that calls for the search and comparison of multiple sources with the aim of arriving at an agreed outcome, i.e. a specific location, hotel, or other item of choice. We felt this stylistic similarity would permit generalisation of insights over different tasks.

Table 1 also displays statistics regarding distinct search sessions, the number of items captured during each task, and total duration of use. We requested these figures from participants after task completion (this helped prevent advance notice of scrutiny impacting behaviour during searches) and all participants consented. It is worth noting that the total duration of use refers to the elapsed time between participants’ first and last search sessions during the entire study. This means that aggregations of completion time for different tasks do not always equate to total duration of participation; some tasks were completed concurrently, whereas others were distinctly separate. Although the lack of precise log data makes these somewhat coarse indicators of usage, it is clear that all participants used their assigned system for at least one week and engaged in multiple information seeking episodes during that time. This in turn increases our confidence in the meaningfulness of participants’ experiences with their assigned system.

Materials & Procedure

Each pair of participants was provided with all necessary software alongside installation instructions and a briefing script introducing the study, tailored according to the system assigned to each pair. The script clarified that participants should use the software we had provided each time they were searching for information towards their chosen task. The script also stated that there were no expectations about the way in which the tool should be used. Rather, participants were encouraged

to use their tool in whatever way they deemed appropriate. We also provided a detailed instruction document explaining the functionality of the relevant system. The document described each system feature in full by providing a screenshot alongside explanatory text.³ All features were given equal descriptive treatment (roughly one page per item) so as not to suggest the importance of any feature over another. We were also careful to ensure that we only described the workings of each feature, rather than how it should be used to support search behaviour.

After installing the relevant system, participants created their own user accounts and were allowed to search at their leisure. As mentioned above, no constraints were placed on participants in terms of process. Participants were free to decide when to engage in search sessions, how long to spend on each session, and when to terminate their information seeking activities. Participants were asked to email the first author once they felt they had satisfied their information need and no longer had use for the system. A semi-structured interview was then performed with each participant. Thirteen of these interviews were conducted face-to-face in a quiet office; two were over the telephone; and one was via Skype. Interviews were one-to-one between participant and the first author; searchers did not, therefore, discuss their information seeking behaviours while their partners were present. This allowed us to cross-check statements and ensure consistency of behavioural accounts. During interviews, participants were invited to access their assigned system if they could not remember exact details or wanted to elaborate on particular functionalities.

All interviews lasted less than one hour. To avoid response bias and to dissuade participants from trying to please us with their answers [4], we began by reminding participants that we did not design the system and there were no right or wrong answers insofar as this study was concerned. We used a basic framework of 30 questions (see Appendix A), which were directed by the concerns of our study and questions used in prior work on collaborative search [3]. These questions allowed us to explore participants’ search process, interactions with their tool and collaborator, and how search products were used in accordance with their chosen tasks. As our protocol was semi-structured, we were able to probe issues as and when they arose, and we invited participants to elaborate on their remarks, allowing novel topics to emerge.

³ All documents are available on request.

Analysis

All interviews were recorded and transcribed by the first author, allowing for early familiarisation with the dataset and development of initial understanding. We used open, axial, and selective coding [22] to identify initial codes, structure codes into concepts, and then group concepts into themes relating to participants' search process and tool usage. The coding process was iterative; transcripts were read several times and initial codes were evaluated and refined during each reading, with internal consistency strengthened by scrutinising the data for counter-examples. We should acknowledge that some codes were influenced by our interview questions. For example, by asking about division of labour, it was inevitable that a related code would arise during analysis. However, many new codes also emerged, indicating to us that the results had novelty beyond current knowledge.

RESULTS

To bring clarity to our results while making sense of our participants' experiences, we cluster our themes using three categories: *search process & management*, *appropriations & afforded behaviours*, and *interface design issues*. The first encapsulates general information about our participants' collaborative search process; such insights are detailed independently of specifics concerning how system features were used. The second pertains to the emergent forms of interaction centred around the use and appropriation of particular system features in support of collaborative search. The final category allows us to drill down into specific issues linked to interface design, many of which are best understood through consideration of participants' search behaviours. As our data are primarily qualitative, we present themes alongside direct quotations from our participants. To identify speakers, we use the form [Px,y] where *x* indicates the ID number of the pair, as listed in Table 1, and *y* refers to the first or second member of that pair.

It is worth noting that participants were generally positive about the design concepts behind the two tools, and all participants believed that their assigned system had merit over ad hoc solutions. Regarding use of system features, all participants stated that they used their system's toolbar, and that they had accessed the shared online space to view search results. In Coagmento, searchers reported using bookmarks, annotations, and snippets. All participants stated that they did not use recommendations, the shared editor, or any of the resources in the sidebar. In Diigo, participants reported using highlights, bookmarks, and screenshots. All participants stated that they did not use Send, Read Later, or any of the functionality within the sidebar. While the absence of log data means that we cannot definitively state that particular features were or were not used during our study, participants' statements are revealing about preferences for some features over others. Before providing more specific details about how features were used to support collaboration, the following subsection paints a general picture of participants' search process and management strategies.

Search Process & Management

Circumstances of Search

Table 2 outlines the search scenarios reported by each pair of participants. Although we cannot specify exactly how many

Pair ID	Co-located		Distributed	
	Synchronous	Asynchronous	Synchronous	Asynchronous
1	✓	✓		✓
2		✓		✓
3		✓		✓
4	✓			✓
5			✓	✓
6	✓			✓
7	✓			✓
8	✓	✓		✓

Table 2. Participants' reported circumstances of search activity. A ✓ in a cell indicates that participants engaged in the relevant type of search.

times each type of search occurred, we did ask participants which of their identified scenarios occurred most often. Responses indicated that distributed, asynchronous search was the most frequent scenario for pairs 1, 2, 3, 5, and 6. Pairs 4, 7, and 8 identified co-located, synchronous search as their most common scenario. While one might expect that such choices would be determined by living arrangements, this was not always the case. Pair 2, for example, were a cohabiting couple who reported conducting all of their searches asynchronously.

Perhaps a more valuable aspect of these results is the fact that tool use was not isolated to a single setting. Rather, all pairs reported engaging in at least two different scenarios of search, and participants described how search activities shifted according to everyday life. Pair 1, for example, reported working only asynchronously when apart, but engaged in both synchronous and asynchronous search sessions when co-located at the weekends. Our participants did not report any immediate difficulties in managing these transitions; this may be due to the fact that, as identified earlier, both tools used in our study are general purpose systems intended to support search across a range of scenarios.

Task Management Strategies

Four pairs reported using division of labour as a way of organising their task. Pairs 1 and 2 divided their travel planning by determining a set of locations and splitting them to create targeted individual searches. Pair 3 established a very distinct role assignment, where different aspects of the search process were divided using semi-formal roles: one searcher was responsible for finding a venue for their band's concert, while another sought places to promote the event. Finally, pair 7 reported dividing the web space during house hunting, with one searcher checking private rental sites and another scouring university-hosted pages. These strategies reflect variations of the *divide-and-conquer* approach described by Morris [14].

The remaining four pairs did not plan a division of labour. These pairs described a 'dive in and do it' approach, akin to the *brute force* management strategy identified by Morris [14]. However, participants did not report any negative experiences related to organising the task in this way.

Search Process

While we were not able to log behavioural data about search activity, participants did provide verbal accounts of their search process, and these are instructive in terms of understanding

what searchers were actually doing during our study. All participants reported following a broadly similar process, beginning with the use of familiar search engines to learn about the information space at a high level. Participants then reported narrowing their search focus to be more specific:

"We started by simply Googling what we're looking for, and when you get leads you search for those specific things, like that specific realtor or a specific area." [P6, 2]

Participants reported that this process was directed by their existing knowledge, and by external factors such as an available shopping budget, dates of travel, or availability of particular venues in the local area:

"I already knew which websites to target, I started off with a few commercial websites and others, looking at StudentPad and the [University] noticeboard." [P6, 1]

After identifying relevant websites, behaviour shifted to the selective capture and bookmarking of results. For all our participants, search and bookmarking were fragmented over multiple sessions. Searchers stated that sessions were interspersed with communication, both during and between searches. As well as ordinary conversation, email, and message exchanges, participants reported using system features like comments and annotations to further their task-related discussions. Participants reported that their overall tasks would end with discussion of the results to reach consensus over an outcome, i.e. the 'chosen products' from their search. In broad outline, this process corresponds well with existing theoretical accounts of web search, where exploratory information seeking has been characterised as an extended, fragmented process, involving multiple search sessions, queries, and information sources [7].

We now delve into the way in which system features were adopted in support of this general pattern of collaborative search behaviour. Our later Discussion section then considers the broader implications of the results, in terms of lessons for collaborative search systems.

Appropriations & Afforded Behaviours

Page Capture for Re-Access and Suggesting Relevance

During their information seeking process, all eight pairs made use of system features intended for information capture: bookmarks, snippets, and annotations in Coagmento; and bookmarks and highlights in Diigo. At the most basic level, these features were used to overcome the ad hoc practice of sharing hyperlinks via email. Participants valued the ability to save pages to a joint repository and, correspondingly, to see pages saved by their collaborator:

"I thought it was really helpful to know exactly what he had looked at and what he wanted to go see. Instead of him having to send me a variety of different emails and links, it's all been saved visually right there for you." [P4, 1]

In providing different features for capturing information, the designers of both systems have anticipated that searchers may wish to retain different types of content. For instance, bookmarks are intended as a way of saving whole pages, whereas

annotations and snippets allow capture of particular subsections or page elements. Our participants, however, tended to adopt just one of these features to achieve the same outcome: that of capturing pages to 'pull out' particular results from the web space. Pair, 1, for example, used Coagmento's 'Snippet' feature to achieve this, whereas pairs 2, 3 and 4 used 'Bookmark' for the same purpose. Furthermore, responses indicated that capture behaviour was more nuanced than simply 'saving pages for later'. First, capture was intended as a way of favouring particular results for more rapid re-access:

"It was particularly useful because I could go back quite easily and know what I'd already put up there, so I wasn't reposting it or covering the same ground." [P5, 2]

Second, capture was sometimes used as way of increasing the salience of particular results, with the aim of suggesting potential relevance to a partner. The intent was not necessarily to return to a particular page but was instead an attempt at bringing something to the partner's attention:

"Looking through Expedia, when I found some good hotels I would snippet them and then I would say to him, oh I've found these ones that look good, and I've snipped them." [P1,1]

Forming and Discussing Shortlists

We identified that page capture was actually related to a more general strategy—that of *shortlisting*, where features like bookmarks and snippets were used to form lists of potential choices and, eventually, reach an agreed outcome. All participants engaged in this behaviour:

"We would aim towards a list of five nice hotels" [P2, 2]

"In the final decision, we had a top five possibles to choose. Which is where the bookmarking came in." [P4,2]

Participants testified that shortlisting was an existing behaviour, but espoused the benefits of the tools in terms of bringing structure to this process:

"Normally we would put links on Facebook, drop all the links into a conversation. But that is chaotic. So it was nice to have them all there without going through our conversation trying to find the link we're talking about." [P7,1]

Participants reported engaging in conversation about their shortlists, either verbally or, in the case of Diigo, by annotating shortlisted links using comments. This occurred as a back and forth process of contributing and evaluating suggestions to the list. Participants stated that using comments on Diigo was useful in supporting this behaviour, indicating the value of features that allow discussion of shortlisted items:

"We would keep commenting on each other's things, so one link had three comments, I would comment and she would comment, then I would comment back. And they weren't all mixed up." [P6,1]

While the actual content of the lists was qualitatively different between tasks (i.e. some pairs were choosing houses, others shopping items) the overall goal was the same: form a list of candidates, discuss and refine the list, and then settle on a 'good enough' outcome in line with the goal of the information

seeking process. However, this process was not always linear or in sequence. Some participants reported that they would be more careful about relevance by evaluating sources on-the-go, keeping only ‘definite’ possibilities in the shortlist. Others would take a more carefree approach, forming a longer list of potentially relevant sources and then narrowing them down at a later stage. Participants reported that shortlists were eventually used and reviewed to reach final consensus:

“A lot of it was down to seeing what the other person thought, and then between the two of us, we decide, oh let’s go for that one. If there’s no availability there, you go to the next choice on your list. So you almost rank them between the two of you, for personal preference.” [P1,2]

Sharing versus Saving: Sensitivity to Time Constraints

In forming shortlists, participants reported that if their task was time pressured and needed to be completed in a short time frame, they would change their behaviour to ensure that their partner checked recorded links more quickly. In other words, rather than wait for their partner to check the system at some unspecified time, participants would use external communication methods to alert their partner about the information:

“We want a house quickly because it’s a little late, so we had to speed up the process. I had to call her up and tell her ok, please go and check this house, I’ve posted it on the group. It’s possible she would go and check, but we were in a hurry, so I have to convey the message to her ASAP.” [P6, 1]

Participants also explained that information was itself often time-critical, in that pages could be subject to expiry. For instance, holiday deals and special offers can be available for limited periods of time, and houses can disappear quickly from the rental market. The knock-on impact for collaboration was one of requiring fast action:

“If it comes up with a great hotel price you click on more details and it would say, there’s only two more rooms with this price left... if there’s only two rooms left, we need to be quick otherwise we might not be able to get this room at this price.” [P1, 2]

However, participants spoke of using out-of-band channels to notify their partner more quickly in this instance:

“Because it gets updated quite often, London seems like quite a quick turnaround on properties. So we chatted over Facebook a bit and then I could send links to websites over Facebook and say does this look ok?” [P2, 1]

This behaviour is interesting given that both systems have features allowing webpages to be sent directly to the email inbox of a collaborator, both of which might be used in a time-pressured search scenario. Yet none of our participants used these features; instead, ad hoc methods of communication were preferred.

Interface Design Issues

In addition to the behavioural patterns identified above, participants described a number of experiences and issues related to interface design and collaborative search.

Shared Search Histories: Information Overload

Participants described how, during search sessions, they would visit the online repository of their assigned system to review their own results and explore what had been found by their partner. Recall that our first system, Coagmento, provides a complete history of search behaviour by capturing all pages visited by the user during information seeking episodes. Previous work suggests that such ‘search trajectories’ can be beneficial for sensemaking and facilitating awareness [18, 19]. However, seven out of eight Coagmento users actually reported that they found the history overwhelming:

“I’d open up the hostel search and there’d be a list of a hundred hostels... and when I’d try and go back to look at which ones were the cheapest, I’d have to look through all of them again to try and find the one in the history. I ended up just searching it again because it was too hard to find.” [P4, 2]

This problem was partly related to the collection of landing pages, portals, and otherwise irrelevant material that held little utility for sensemaking. Participants suggested that they wanted to get rid of such pages and see only the most relevant information without having to trawl through the entire history:

“I never looked at the history because the snippets were the core information. I just want something that jumps out and says... this is the information that she has found.” [P1, 2]

Participants suggested that it would be helpful if Coagmento’s history was searchable, or that, rather than allowing the system to indiscriminately record all pages, the process of information capture could be more selective:

“It’s almost like when you’re searching you want to be able to dictate what’s saved. So even if it’s just a little ‘plus’ button, where maybe you don’t want your Google search to come up, but, whatever site, even if you don’t think it has anything to do with it, but you might want to look at it again, you can just hit a button and it’s going to come up there.” [P4, 2]

Some users did mention benefits of the timeline view. One member of pair 1 reported retracing their partner’s search results to see what had been done earlier. Members of pair 2 reported using the timeline to direct their own work based on the results of their partner, while one member of pair 4 thought the ability to rediscover past results could be useful.

Responses from Diigo users provide an interesting contrast on this issue. Such participants did not express any dissatisfaction about not being able to view a complete search history, and none expressed any desire to have more information about their partner’s search process. Some were actually skeptical about the need to understand their partner’s process:

“As long as you both know what you’re looking for, it doesn’t matter how you go about finding it. I didn’t feel like I needed to know what he typed into Google.” [P8, 1]

Shared Search Histories: Privacy Concerns

The fact that Coagmento captures images of all pages visited made three users hesitant about the potential consequences of having personal information recorded:

“While I’m searching for hostels I’m using Facebook and Gmail

at the same time... I think if you're working in a group that history could get a little weird, because you're going back to your bank account to figure out how much you still have left." [P4, 2]

This highlights how, for our participants, real-world collaborative search was not an isolated endeavour; instead, it occurred as part of a broader planning activity that required access to personal and private information. Clearly a screenshot image of such content would risk a privacy violation. Participants were not specific about whether any such violations occurred during the study, but the mere threat of information leakage was enough to affect behaviour. In line with earlier statements concerning information overload, participants suggested the solution of selective tracking:

"I was always making sure that I wasn't logging in to other things while logged in to that. Maybe if it had a way of knowing it was on a page for logging in to your emails... it would know not to record those kinds of pages." [P3, 2]

In contrast, privacy did not arise as an issue during discussion with Diigo users, likely due to the fact that no pages are captured unless specifically requested by the user.

Shared Search Histories: Sensemaking and Rationale

Several issues arose related to sensemaking of the shared history, in terms of understanding *what* and *why* information had been found. Regarding the former, responses suggest that searchers need to be able to appraise page representations both quickly and easily, and do not want to spend a long time making sense of shared results. This was exemplified by participants' experiences with Coagmento, where each item in the shared history appears as a small thumbnail showing an image of the webpage (see Figure 1). However, rather than simple thumbnail images, participants wanted an up-front summary of contextual information about each page: a descriptive headline, a preview of its contents, and the time it was saved. We note that this contextual information is available within Coagmento but requires the user to click on each thumbnail to view it. Our responses imply that such information should be presented up-front, with low interaction costs, to facilitate rapid assessment of relevance and sensemaking of pages:

I would want to see the header you'd get if you bookmarked it in a browser. So, under 'hotel search', then the name of the hotel, that sort of thing. The thing is that it's either two or three links you have to click through before you even get to the page itself... And so I would've used that more if I'd had the ability to go, I know what that is... that's what I want... without having to click, click, click." [P2, 2]

Although Diigo records pages using a title and contextual information, sites are not listed with thumbnail images (see Figure 2). Somewhat ironically, all Diigo users indicated a desire for thumbnail previews of each page, as with the style used by Coagmento. This suggests providing both types of information could be beneficial for sensemaking:

"You know like on Facebook when you put a link they give you a small image and a description of what's in the website... it would be nice to have that feature because it triggers your

mind about the thing you saw... a picture of the house could appear there automatically." [P7, 1]

A second issue related to sensemaking was that searchers wanted to understand *why* particular pages had been visited by their partners. This was true of both the search history and pages shortlisted using bookmarks:

Without the other person telling you what they'd gleaned from each of the links, it was difficult. You need some explanation, he could have looked at all this stuff and thought, this is a load of rubbish. Then what's the point of you looking at it? [P1, 1]

Participants wanted to annotate and append specific pages to provide rationale about why results had been selected, and to draw attention to specific aspects of webpages. Both practices are in line with the act of suggesting relevance:

"When I would bookmark a whole page, it wasn't necessarily the whole page that I wanted her to look at... I wanted her to pick out the Regent's Canal, but if I'd bookmarked the page, she wouldn't necessarily know that." [P4, 2]

Users of Diigo reported using a short description at the point of page capture to inform their partners about why particular links were relevant. This information would then appear alongside each result in the shared space. Searchers also described appending links with information gleaned from the page so that their partner did not necessarily have to re-run the search. However, users also mentioned that they wanted more ways of tagging or marking captured results. One goal was to indicate that certain pages were temporarily irrelevant. Users did not want to delete such pages in case they became useful later on:

"I added 'TAKEN' to the house because I didn't want to delete it, it's better to have it there for future reference because you never know, so I just put that tag." [P7, 2]

Segregation and Manipulation of Shared Information

Our chosen systems allowed shared results to be viewed in two forms: either as a list of one's own results, or as a combined list of with those of one's partner. However, four participants desired to see *only the results of their partner*. This action was not possible in either system, but this type of separation would clearly be advantageous when attempting to make sense of a partner's work:

"It came up with all of your stuff to start with, what you'd done when you go on it... which I really didn't care about as much, because you know what you've done. That's not the immediate reason you go on there, you go on to see what the other person has done." [P1, 1]

Participants also wanted the ability to distinguish different aspects of their work from others. For example, pages related to flights, hotels, and sightseeing were treated equally by the systems, i.e. as superficially equivalent 'pages'. But to searchers, these pages were related to different subtasks. Users wanted more ways to differentiate and classify these results by creating subcategories and folders that would allow results to be directed towards particular sections. Only one of our tools, Diigo, provided support for this. Participants were able to create specific groups to manage the overlap in

concurrent completion of different search tasks:

“We made different groups, one was for house hunting, one was for holiday hunting, and the third one was on these dresses that we were shopping for online.” [P5, 1]

Additionally, users reported that they wanted the ability to manipulate the results in greater depth; for example, by moving sources from project to project in accordance with changes in information need and the overall progress of task completion. Users of both systems also wanted to rank and reorder results in their shortlists, so as to allow for visual comparisons. This related to the practice of shortlisting and reaching consensus:

“if I have a preferred order for all these houses, like if I want to call this one first and that one second, I would like to have the possibility of rearranging them.” [P7, 2]

Awareness and Notifications

In discussing their partner’s activities, participants’ responses imply that notifications of recent work (e.g. searches, bookmarks) would help to maintain a general level of awareness:

I think it would be nice to get actual notifications of things. If she’d bookmarked a page, I could get a notification of a bookmark, or a tag, something like that. [P4, 2]

Interestingly, both systems do provide notifications but in different ways. Coagmento’s notifications appear in the sidebar but, as all of our participants stated they did not use this feature, likely went unnoticed. By contrast, Diigo provides daily email notifications alerting collaborators to the presence of new links in shared groups. Although one participant found the emails useful in terms of being able to check recent links with her mobile device, the remaining participants said they did not actually read the emails; instead, they were mainly used to gain awareness of the mere fact that *something* had been done. Again, this speaks to a desire for some general level of awareness about the fact that contributions are being made. When checking their assigned system, participants wanted to be able to find these contributions quickly and easily:

“It could highlight that this is the new comment, or this is the new thing, instead of me having to search constantly.” [P6, 2]

Two users stated that it would be helpful to have awareness of when their saved results had been viewed by their partner, i.e. a confirmation that results had actually been seen. Responses also indicated a desire for such notifications to be more immediate and situated in the web browser, rather than the system itself. Several drew on their experiences with Facebook to suggest how these notifications might be implemented:

“It could work like Facebook, where you get the notification, you click on it and it takes you to the page.” [P4, 2]

Effort Requirements

Finally, throughout our dataset, there was a general undercurrent of wanting a minimal threshold for effort. The perception was that, when search was the primary task, any additional effort above and beyond an ad hoc solution was undesirable:

“All you wanted was something easy, you didn’t want something that was going to add a load of time, because the search was

quite time intensive. You didn’t wanna feel like you were going to have to do loads of extra work on it.” [P1, 1]

Users of Coagmento made more specific remarks about effort, in that they would prefer aspects of the tool to have lower interaction costs. In addition to the earlier statements concerning the number of operations required to access contextual information in the search history, participants remarked upon the excessive number of stages involved in other tasks. For example, when the system asks users to rate and leave comments about bookmarks at the point of page capture:

“I guess the reason the bookmark thing asks for a one to five rating is so it could do some filtering. But at the time it’s too much effort. I don’t want to do that, I want to automatically know what’s important.” [P2, 1]

DISCUSSION

The present study sought to investigate the use of collaborative search tools in everyday settings. At a broad level, the fact that searchers found merit in their assigned system suggests that the design concepts embodied by each are on the right track—it was not the case that participants considered the tools unusable or inappropriate, and many stated informally that they planned to use the tools again outside of our study. Some of the primary benefits noted by our participants included the ability to save links to a shared space, thereby negating the need for ad hoc workarounds, and the fact that the tools brought structure and persistence to the otherwise ephemeral process of collaboration over time.

General aspects of participants’ search behaviour are revealing about how the tools were used in everyday circumstances. Beginning with the search process, collaboration among our participants was not isolated to a single configuration of time and space. Instead, all participants reported engaging in multiple scenarios. Although no work has argued otherwise, these findings suggest that tools should aim to support search across a range of settings rather than a single time/space configuration. The systems we used were not restrictive in this regard, but some other tools (e.g. [2, 13]) have been designed with specific circumstances in mind.

One of our most consistent behavioural findings was the generalised use of page capture features (snippets, bookmarks, and annotations). While this stresses the importance of allowing searchers to save pages to a shared space, our results suggest that capturing pages is actually more nuanced than simply ‘saving pages for later’. Rather, searchers used these features to save pages for re-access; to share results with their partner with the aim of suggesting relevance; and to form *shortlists*, i.e. subsets of results that are collected as potential candidate outcomes for the group’s collaboration. However, we saw that pairs tended to settle on, and persist with, a single means of page capture to achieve all three (e.g. Pair 1 used only snippets for shortlisting and suggesting relevance). Given that these are essentially distinct behaviours, which might appeal to different users in the light of different nuanced needs, we suggest that each is potentially worthy of independent support. Future tools could provide separate methods for *page saving*, *page sharing*, and *suggesting relevance*. Regarding the latter,

some existing tools allow collaborators to suggest relevance by ‘liking’ pages [7] or through up- and down-voting of results [2], but future systems could do more to separate shortlisted pages from those that are awaiting judgements of relevance. This would avoid conflation of the two, as occurred for some of the participants in our study.

In forming shortlists, participants wanted to annotate saved results to share rationale about *why* pages had been selected; to make pertinent information salient (e.g. the rent of a saved property or the price of a hotel); and to engage in task-related discussion about the quality or relevance of results. Several systems [16, 18] allow commenting of saved items in the manner offered by Diigo; our study reiterates the value of allowing annotation of links to facilitate sensemaking and sharing of knowledge in the context of everyday tasks.

However, we found that searchers wanted to rearrange listed items to express an order of preference, or ‘check off’ items by deleting some while retaining others. This process of bringing structure to data is part of sensemaking [23] but our chosen systems did not allow users to repurpose results in any meaningful fashion at the end of their tasks. The implication here is that, as well as supporting aspects of the search process, future tools might also do more to support the manipulation and reuse of search *products*. Since searchers want more ways of interacting with their results, different workspaces could be provided to allow segregation of items; searchers could then be allowed to move items between these spaces in accordance with relevance judgements and task progress. Such an approach can be seen in the ViGOR system [13], which includes a workspace that allows searchers to drag and drop results into groups that can then be reorganised and restructured at will.

Some broader issues related to sensemaking and awareness concerned searchers’ difficulty in understanding *what* had been found during their partner’s searches. This is a foremost challenge during collaborative search [18] and indeed this was true for our participants. An initial issue concerned identifying the location of recent work when checking their assigned system—searchers often found it hard to identify recent collaborative contributions (e.g. new annotations) in the presence of a large search history. Future tools can benefit by ensuring that notifications are available and that they lead straight to recent contributions. A second lesson is that sufficient information must be given to allow searchers to appraise shared representations with relative ease. Prior laboratory studies suggest that rapid sensemaking can be promoted by displaying contextual information about search results [18, 19]. Our participants’ statements align with this, and indicate that representations should combine content previews (e.g. a visual thumbnail) with contextual information (e.g. a page description) that allow comprehension while lessening the need for collaborators to click through and revisit every individual result.

An important design challenge raised by the present study concerns the vexed question of history logging and the extent to which page capture should be deliberate or automatic. Earlier work suggests that the approach of capturing all pages, which provides a persistent account of a collaborator’s ‘search trajectory’, is beneficial for sensemaking [18, 19]. However,

we found that, with this approach, the quantity of information captured was regarded as overwhelming. Rather than wade through a large history, our searchers simply wanted to identify only the most relevant information. This finding is similar to that obtained in a laboratory evaluation of CoSense [19], where searchers reported feeling overwhelmed by search histories and struggled to identify ‘good’ from ‘bad’ information. In suggesting solutions to this problem, Paul & Morris [19] recommended that designers allow filtering of search histories by content. This solution would not, however, necessarily overcome the problem of information overload as pages accumulate over time—even filtered lists could become difficult to comprehend. One interesting aspect of our results is that none of our Diigo users felt the need for a complete search history. This tentatively suggests that a complete history may not always be necessary for quotidian tasks, especially if collaborators communicate about their progress as part of their daily routines. (As in our study.) That Diigo users did not complain about the effort of saving pages also hints that page capture could be left entirely under the control of the user. This would, in turn, prevent the collection of interstitial pages that appear to offer little utility for sensemaking.

This work has also identified issues that arise only in real-world search. One example pertained to time-criticality, both of search processes and information itself, with some items available only for short windows. We are not aware of any prior work that has raised time constraint as an issue for collaborative search—perhaps future systems could elaborate on this issue and provide more pointed ways of delivering results to one’s partner (e.g. through mobile devices). A second issue was the subject of privacy, which highlights the need to ensure that tools are capable of interleaving with the broader tasks in which collaborative searches are embedded. Searchers in our study were nervous about the potential capture of personal or private material. While an obvious answer to this issue would be to allow deletion of pages from the system’s history, a more elegant solution would be to record only pages relevant to the user’s information need. This could be achieved by allowing the user to enable or suppress tracking in specific browser tabs—this would compartmentalise search to a specific location, leaving other areas safe for private multitasking.

As many of these implications concern specific issues, it is worth considering how the present study can direct future collaborative search tools more generally. First, our results lend support to Morris’ recent assertion that collaborative search solutions must be low-effort and “sufficiently lightweight compared with status quo ad hoc solutions” [15]. Our participants expressed dissatisfaction if interactional demands imposed by their tool exceeded those of their previous solution (e.g. email, Facebook). Since the implication here is that any tool with unnecessary effort requirements stands to fail, features in collaborative search systems should be benchmarked against the equivalent ad hoc solution. Additionally, the tools used in our study included many features which, according to our participants, were not necessary for their task. This suggests that future solutions could be scaled back in favour of lightweight support for core collaborative search behaviours. An example system might support the rapid sharing of pages between

two linked browsers, with simple awareness mechanisms, like those desired by our participants, that notify collaborators of recent activity after their web browser is opened.

Alternatively, more specialised systems could be developed to provide targeted support for specific tasks. Some of our participants mentioned informally that they would value support beyond that for search. In travel planning, for example, searchers could be able to apply their shortlists to a map, helping them with their journey or sightseeing plans. The wider implication here is that collaborative search tools could be embedded in larger applications that support a broader range of high-level planning tasks.

Limitations & Future Work

One limitation of this study is that we were only able to report on a small number of pairs from a single culture. Larger groups of three or more could differ in terms of their search and sharing practices. Since demographic, cultural, or other socio-economic factors may also impact search behaviour and system usage, future work should explore these issues.

While our study results were based on search tasks with high external validity, the patterns of behaviour reported here may not generalise to every other form of collaborative search. For example, we were not able to examine the very commonplace task of literature review, an endeavour that will no doubt be familiar to readers of this paper. Future work should examine how our findings map to other tasks.

Finally, this work is limited in ways that would be encountered in any study involving retrospective self-reports. Our results are reliant upon what our participants were able to recall about their experiences, and might also have been tempered by our choice of interview questions. We were unable to access quantitative data collected at the system level—future work might employ such an approach to reach more precise measures of tool use. Reliable data about which features are, and are not, used in field settings would help to confirm our findings and guide the design of collaborative search tools.

REFERENCES

1. Amershi, S., and Morris, M. R. Cosearch: a system for co-located collaborative web search. In *Proc. CHI* (2008), 1647–1656.
2. Capra, R., Chen, A. T., Hawthorne, K., and Arguello, J. ResultsSpace: An experimental collaborative search environment. In *Proc. ASIST*, vol. 49 (2012), 1–4.
3. Capra, R., Marchionini, G., Velasco-Martin, J., and Muller, K. Tools-at-hand and learning in multi-session, collaborative search. In *Proc. CHI* (2010), 951–960.
4. Dell, N., Vaidyanathan, V., Medhi, I., Cutrell, E., and Thies, W. “Yours is better!”: Participant response bias in HCI. In *Proc., CHI ’12* (2012), 1321–1330.
5. Foley, C., and Smeaton, A. Division of labour and sharing of knowledge for synchronous collaborative information retrieval. *Information Processing & Management* 46 (2010), 762–772.
6. Golovchinsky, G., Adcock, J., Pickens, J., Pernilla, Q., and Back, M. Cerchiamo: a collaborative exploratory search tool. In *Proc. CSCW*, Demo Session (2008).
7. Golovchinsky, G., Diriyee, A., and Dunnigan, T. The future is in the past: designing for exploratory search. In *Proc. IIIIX* (2012), 52–61.
8. Golovchinsky, G., Diriyee, A., and Pickens, J. Designing for collaboration in information seeking. In *Proc. HCIR* (2011).
9. Golovchinsky, G., Pickens, J., and Back, M. A taxonomy of collaboration in online information seeking. In *Proc. JCDL Workshop on Collaborative Exploratory Search*. Pittsburgh, PA. (2008).
10. González-Ibáñez, R., Haseki, M., and Shah, C. Time and space in collaborative information seeking. In *Proc. ASIST* (2012), 1–10.
11. González-Ibáñez, R., Haseki, M., and Shah, C. An analysis of communication and performance in collaborative information seeking. *Information Processing & Management* 49 (2013), 1165–1179.
12. González-Ibáñez, R., and Shah, C. Coagmento: A system for supporting collaborative information seeking. In *Proc. ASIST* (2011), 1–4.
13. Halvey, M., Vallet, D., Hannah, D., Feng, Y., and Jose, J. M. An asynchronous collaborative search system for online video search. *Information Processing & Management* 46, 6 (2010), 733–748.
14. Morris, M. R. A survey of collaborative web search practices. In *Proc. CHI ’08* (2008), 1657–1660.
15. Morris, M. R. Collaborative search revisited. In *Proc. CSCW* (2013), 1181–1192.
16. Morris, M. R., and Horvitz, E. Searchtogether: an interface for collaborative web search. In *Proc. UIST* (2007), 3–12.
17. Morris, M. R., Lombardo, J., and Wigdor, D. Wesearch: supporting collaborative search and sensemaking on a tabletop display. In *Proc. CSCW* (2010), 401–410.
18. Paul, S. A., and Morris, M. R. Cosense: enhancing sensemaking for collaborative web search. In *Proc. CHI* (2009), 1771–1780.
19. Paul, S. A., and Morris, M. R. Sensemaking in collaborative web search. *Human-Computer Interaction* 26, 1-2 (2010), 72–122.
20. Shah, C., and González-Ibáñez, R. Spatial context in collaborative information seeking. *Journal of Information Science* 38, 4 (2010), 33–349.
21. Shah, C., and Marchionini, G. Awareness in collaborative information seeking. *Journal of ASIST* 61, 10 (2010), 1970–86.
22. Strauss, A., and Corbin, J. M. *Basics of Qualitative Research*. Sage, 1998.
23. Tao, Y., and Tombros, A. An exploratory study of sensemaking in collaborative information seeking. In *AIRS*. Springer, 2013, 26–37.

APPENDIX A: INTERVIEW MATERIALS

The following script was read to each participant prior to the commencement of each interview.

Thank you for participating in our study, today I will be asking questions about your use of Coagmento/Diigo. Before beginning, I would like to tell you a few things about the interview.

- First of all, please feel free to be frank and honest when talking about your experiences with the system. We didn't build the tool and we have no reason to be offended by anything you say about it, so please feel free to speak truthfully and be honest about your experiences.
- Also, there are no right or wrong answers in this study. So please feel free to share your opinions and ideas, even if you think they might not be relevant.
- I have basic framework of 30 questions but sometimes they get answered during conversation, so I may not need to ask all of them. The interview shouldn't last more than an hour in total. We can follow up on your experiences if there is something particular that you want to talk more about.

Do you have any questions before we begin?

Interview Questions

Overview of the Project

I'd like to start by asking you some general questions about your search activity.

1. What type of information were you and your partner looking for? (*If more than one task: ask for each.*)
2. What was the overall goal of the search? (*If more than one task: ask for each.*)
3. Where did you do most of the searches? (*Where did they take place.*)
4. Did you search at the same time as your partner or at different times? Were you together or apart?
5. How long did it take to complete the task?
6. How many search sessions were there? (*Ask then verify if necessary by inviting them to check the system.*)
7. How would you normally go about looking for information when doing tasks like the one you did during the study?

Information Seeking

Let's talk about the process you followed when searching for information.

8. Can you give me a general idea of your approach towards achieving your goal?
9. How did you organise the task between the two of you?
10. What challenges, if any, did you encounter related to searching and managing results found?

11. What did you do with information once you had found it? (*Probe why they did what they did with the information*)
12. Did you encounter any problems during the execution of the task?
13. How did you resume your searches from previous sessions?

Communication

Let's talk about how and your partner communicated and exchanged information regarding your chosen task.

14. How did you share information with your partner?
15. How did they share information with you?
16. How did you and your partner communicate during the time you were working on your tasks? (*What methods did they use to discuss what they had found?*)
17. Did you ever communicate during searches?
18. What did you do to understand what your partner had searched for and looked at during their work?

Tool Use

Let's talk about the tool you were using.

19. Thinking back to your use of the toolbar, which features did you use most frequently? (*Ask them to elaborate on what they used them for. Ask them about each feature and whether or not they used them.*)
20. Were there any features you chose not to use? (*If so, why didn't they use them?*)
21. Did you use any of the toolbar's communication features, for example the chat or sidebar?
22. In what ways was this tool useful during your task?
23. Were there any problems during your use of the system? (*If so, what were they?*)
24. Are there any ways in which you would improve this tool?
25. Did you use any other tools or methods to capture information? (*e.g. paper notes, spreadsheets*)
26. Can you give me an opinion of what you thought about the toolbar, overall?

Ending the Process

27. How did you decide that the quantity of information you found was enough? (*As in, at what point did they decide to terminate information seeking activities*)
28. How did you achieve consensus regarding the outcome of the work? (*As in, how did they decide which option to select*)
29. Overall, how would you describe your success in achieving what you wanted to achieve by using the tool?
30. Is there anything else you would like to add that we have not covered?