Division of Labour and Sharing of Knowledge for Synchronous Collaborative Information Retrieval

Colum Foley and Alan F. Smeaton

Clarity: Centre for Sensor Web Technologies
Dublin City University
Dublin 9
Ireland

Abstract

Synchronous Collaborative Information Retrieval (SCIR) is concerned with supporting two or more users who search together at the same time in order to satisfy a shared information need. SCIR systems represent a paradigmatic shift in the way we view information retrieval, moving from an individual to a group process and as such the development of novel IR techniques is needed to support this. In this article we present what we believe are two key concepts for the development of effective SCIR namely *division of labour* (DoL) and *sharing of knowledge* (SoK). Together these concepts enable coordinated SCIR such that redundancy across group members is reduced whilst enabling each group member to benefit from the discoveries of their collaborators. In this article we outline techniques from state-of-the-art SCIR systems which support these two concepts, primarily through the provision of awareness widgets. We then outline some of our own work into system-mediated techniques for division of labour and sharing of knowledge in SCIR. Finally we conclude with a discussion on some possible future trends for these two coordination techniques.

Key words: information retrieval, collaborative information retrieval, collaborative search

1. Introduction

Collaborative information retrieval refers to systems and technologies which support collaboration between people during the information retrieval

(IR) process. This research domain is placed at the convergence of the information retrieval and computer-supported cooperative work (CSCW) fields. CSCW systems are often categorised based on their position in the two dimensional space of time and place [1]. Using this taxonomy, most of the work to date in collaborative information retrieval refers to systems and techniques which support an asynchronous, passive, implicit form of collaboration between remote users. The aim of these asynchronous collaborative IR systems is to improve the IR process for an individual by using the past experiences of others. Collaborative filtering systems [2] attempt to provide serendipitous recommendations of items (documents, books, DVDs etc.) to users based on the opinions of others, where opinion is expressed as implicit or explicit ratings. Collaborative re-ranking systems use click-through data associated with previous users' queries in order to re-rank a ranked list for a user [3]. Collaborative footprinting [4] systems attempt to contextualise a user's search process by showing the *trails* left by previous users through the information space, such as showing a popularity icon indicating those links most often followed, or the comments associated with results on a ranked list. Recently with the ever growing emphasis on the *sociable* web we have begun to see the emergence of commercial collaborative or *social* search systems on the web. Collaborative bookmarking websites such as Del.icio.us [5] allow users to tag and share webpages with others. Heystaks [6] enables groups of users to arrange similar searches into *staks* so that group members can see those pages found by their friends.

Although these systems are termed "collaborative", quite often the users involved are unaware that they are collaborating whatsoever. By and large the users of these systems do not know the people they are collaborating with. In essence, the goal of these systems is equivalent to any traditional IR system - to provide effective retrieval of information for an *individual*.

In our work we are interested in exploring the potential for synchronous, explicit, real-time collaborative search in which two or more users cometogether in a focussed group, in either a co-located or remote setting, in order to satisfy a shared information need. We refer to these systems and technologies as Synchronous Collaborative Information Retrieval (SCIR) to emphasise the nature of the collaboration in this setting. SCIR systems represent a significant paradigmatic shift in how we view information retrieval, from an individual to a *group* process. As a result, how we approach IR in terms of information needs, notions of relevance, and ranking for example needs to be re-visited in this setting.

Despite claims that IR has always been a social process [7, 8, 9] and several studies showing that SCIR is common among groups of people in either a work or social scenario [10, 11, 12, 13, 14], SCIR has received little attention from the broader IR community to-date.

In [15] we proposed the concepts of division of labour (DoL) and sharing of knowledge (SoK) for SCIR. Both DoL and SoK have long been recognised as important items in successful CSCW systems [1, 11, 16, 17]. In our work we are interested in exploring their application to an information retrieval environment involving two or more users. DoL and SoK operating in such an SCIR environment have the potential to allow for a more effective group search by reducing the redundancy across collaborating searchers whilst enabling users to benefit from the experiences and discoveries of their collaborators.

The majority of published work in SCIR has focussed on improving the *awareness* across groups of remote collaborating searchers through the provision of various awareness cues. This is not surprising considering that awareness has been *the* central issue in designing systems to support CSCW for many years [18, 19, 7, 20, 21, 22]. By providing users with a perception of their collaborator's activities during the search, the idea is that the users themselves can then co-ordinate the search task. An important design consideration when designing CSCW systems, however, is the trade-off between providing adequate awareness and disrupting individuals work [18]. Indeed as noted by Adcock et al. [23] requiring users to both search and coordinate the group activity may cause searchers to suffer from cognitive overload.

In this article we will look at how DoL and SoK can be realised in SCIR systems. Firstly, in section 2 we will describe field studies which have observed SCIR occurring *in the wild* and highlight how DoL and SoK occurs often despite the lack of provision of tools to enable effective collaboration. In section 3 we will outline how a DoL and SoK can be realised in systems to support synchronous collaborative search, here we will draw on previous work from the general CSCW literature before outlining what they mean for the novel domain of SCIR. In section 4 we will outline attempts made to enable a user-driven coordination of SCIR via the provision of awareness widgets, and show how such techniques can enable a DoL and SoK. Following that, in section 5 we will outline our own work in implementing *system-mediated* techniques to support DoL and SoK in SCIR, where we attempt to move the burden for coordination of the search from the users to the back-end search system in order to reduce the cognitive load on collaborating searchers. We

3

will present our implementation of system-mediated techniques to realise both DoL and SoK for search. We will conclude the article with a discussion on some possible future trends for both concepts.

2. Information Retrieval: A Social Process

Researchers claim that IR has always been a social process [7, 8, 9] and studies have shown how collaboration is commonplace throughout the information retrieval process, despite little provision being made for collaborative tools within IR systems [24]. The collaborative nature of the information retrieval process has been investigated in a number of domains including: academia, industry, medicine, and the military [13].

Twidale and Nichols [10] observed collaboration in activities around 11 OPAC computer terminals in a university library. In the study they found that 10% of all uses of the terminals were of a collaborative nature including: multiple users working around the same terminal in a problem solving task and pointing and gesturing at the screen, users working on adjacent terminals and coordinating their actions, and users asking others questions such as "How do you do that?". They noted that this rich collaboration occurs despite the fact that talking in the library was discouraged and that terminals and information systems were designed for single usage.

In a study of information sharing behaviour across four academic disciplines, Talja [25] demonstrated that collaborative information retrieval is as commonplace as individual or solo information retrieval.

In a long term study of the collaborative information retrieval activities in design teams at Boeing and Microsoft, Poltrock et al [11] observed collaboration occurring at all stages in the information retrieval process (identifying an information need, query formulation, retrieving information). Based on interviews, observations and records of meetings and emails, they concluded that a division of labour was vital to the coordination of the team's work, enabling individuals to perform parts of the work alone, and that a shared information space, where users could share ideas and ascertain information, was a central element of any collaboration system.

In a study of the activities involved in resolving a patent application, Hansen and Järvelin [12] observed that collaboration amongst patent engineers was common in most phases of the patent task with information seeking the most common. They noted that overlap was common across

4

patent searches and that this led to sharing of information such as documents, queries, and opinions. Hansen and Järvelin also observed patent engineers cooperating on work tasks, dividing the task amongst colleagues and sharing search strategies.

In a survey of web search usage amongst workers at a large corporation, Morris [24] found that collaboration was commonplace in web search despite it not being supported in search systems. In the survey, over 97% of all users reported having used some form of collaboration when searching the web. This included: 87.7% of people saying that they had looked over someone's shoulder while they searched and suggested query terms, 30.4% of people who said that they had used instant messaging (IM) to coordinate a real-time search, and 23.5% who had used a large display to perform a web search during a group meeting. For typical group size for collaboration: 80.7% of respondents had collaborated in groups of two, with 19.3% reporting a group size of three or four (no larger group sizes were reported). When multiple users collaboratively searched an information space together using separate devices, Morris [24] identified two common task division strategies: (1.) *Divide and conquer*, where the coordination of the task involved assigning separate subtasks to each individual, which was used by 19.3% of respondents, and (2.) the *brute-force* approach where no coordination took place and users searched separately was used by 24.8% of respondents. The common types of task for which collaboration took place were purchasing items, booking travel, and researching property.

From these different studies, and others reported in the literature [13], we can conclude that collaboration is common in information retrieval, this collaboration can occur at any stage of the IR process, and that collaboration occurs *in spite of* a lack of support for these activities in most IR systems.

3. Division of Labour and Sharing of Knowledge for SCIR

Both the division of labour and the sharing of knowledge have long been identified as important aspects of any successful CSCW system [16, 17, 11, 1, 26]. Both concepts form part of the broader *coordination theory* for CSCW [27], a set of principles which define how collaborative work can be implemented such that collaborators can work together harmoniously.

3.1. Division of Labour

Division of labour refers to the process of dividing up the group task across collaborators in order to share the workload across the group. An effective DoL policy should enable the group as a whole to work more efficiently by reducing the redundancy across collaborators thus ensuring that no two users spend time completing the same task unnecessarily.

Some form of DoL can be seen across most examples of CSCW applications in the literature. Furthermore the division of labour across workers has long been studied in workplace. In [28] Poltrock and Handel used both ostensive and performative modelling techniques in order to examine the subdivision of processing in organisations such that it could inform the development of collaborative technologies.

From examining the CSCW literature we find several ways in which a DoL can be achieved. In [26] Sharples describes three strategies for collaborative writing: *parallel*, where the work is divided into subtasks and all collaborators work simultaneously; *sequential*, where a task is divided into stages such that the output from one stage is handed onto the next user; and *reciprocal*, where group members work together mutually adjusting their activities. Roles have also been investigated as a means of dividing work across users working together, where each user assumes responsibility for certain jobs according to their role. Within an organisation, the role of a manager is often to divide the activities needed to complete a project amongst a team of workers, where the roles of individuals in the organisation in terms of their skills and expertise determine the job they are assigned [11]. *Protocols* either technological or social can also be used to divide the task [1]. A technological protocol in a co-authoring CSCW application could, for example, lock a section of a document that was being edited by one user such that other users could not change it. Social protocols do not enforce any system-level restrictions but instead rely on etiquette and user-user coordination. The Grove collaborative writing system proposed by Ellis et al. in [1] operated such a social protocol, in which collaborating users had free reign to edit any part of a document at any time. The result was not chaotic but rather, after a learning period, the social protocol was found to be successful at mediating collaboration.

It is clear that a strict division of broader tasks into subtasks to be allocated to each user is more suitable when there is a clear boundaries between the different stages of the task. An information retrieval process is by its nature a very flexible and iterative process: users can rapidly query, view results and examine the contents of documents. As a result the boundaries

between the different stages can become difficult to distinguish. Implementing a stage-by-stage division of the task is therefore difficult in modern IR systems [29]. As a result, as we will see in section 4, the vast majority of work to date in SCIR have relied on social-protocols and awareness cues as a means to divide the search task. A DoL in SCIR means dividing the search task in such a way that each member of a collaborating group completes a piece of the search task. When users search to satisfy the same information need, as is the case in SCIR, they often use the same or very similar query terms [30]. If searching within the same corpus of documents, this similarity across queries will result in highly similar ranked lists being returned to users. Therefore, without some form of DoL policy in SCIR, users may spend a lot of time replicating the work of their collaborators. There are numerous ways in which a DoL can be achieved in SCIR. For example each user may explore a subset of the information space during the search, either by searching through different collections or searching using different query terms for example. Users may assign themselves different parts of the search task so that one user performs query formulation and another relevance judging. If done effectively, DoL should enable the group to gain a greater coverage of the information space by reducing the redundancy across the group and allowing a greater throughput of document assessment.

3.2. Sharing of Knowledge

Sharing of knowledge refers to the passing of ideas and information between collaborators during a group activity. The ability to effectively share information has been recognised as one of the foundations of any group activity [17]. In any group work setting each user will bring their own experiences and expertise to the shared task. Effective SoK policies should allow collaborators to exchange information and ideas seamlessly across the group to enable users to benefit from the expertise of their collaborators.

CSCW applications have attempted to support the sharing of ideas and knowledge through the provision of shared workspaces [19, 11, 31, 32, 33]. Workflows in organisations such as those modelled in [28] have implicit knowledge sharing between the different stages in the processes as one worker or department passes information on in the chain. Informal and spontaneous sharing of knowledge between knowledge workers is seen as critical to the way many organisations operate [34, 11, 19]. Knowledge sharing is particularly evident across the intensive patent application domain [12]. Large multinationals often have internal systems whereby potential patents applications

7

can be posted on a shared workgroup system for critiquing by colleagues [35]. Such a system can enable fellow patent engineers to discuss the application and to provide information on conflicting patent applications which they may have uncovered in the course of their work. Given the high-cost associated with patent applications, such a corporate-wide knowledge sharing is extremely important in the patent domain.

Gutwin et al. provide an overview of shared workspace technologies in [19].The advantage of using an online medium for communicating ideas and discussions is that it can be accessible across the entire organisation, even to those whom it was not originally intended. Activity Explorer [36] allow a shared access to repositories of documents or artifacts in organisations. People can access the shared repository, manipulate the shared object and communicate with each other via an instant messaging style communication where they can chat and send snapshots of their work. Lai et al [33] describe a collaborative spreadsheet based system for sharing knowledge between collaborators by sending messages and transferring files. Notecards [37] was an early hyper-text based system for idea generation and structuring. A notecard was an electronic hypertext placeholder containing text and images. Users could share information and critique notecards using annotations. Notecards could be passed between users and each notecard had a history associated with it which contained details of the changes made by participants.

In our work we are interested in exploring the application of sharing of knowledge policies in the context of synchronous collaborative information retrieval. When people come together to search for information, each user will have different knowledge and expertise, for instance, some may be more familiar with the search topic, others may be more familiar with the search tools. Based on their knowledge, users may approach resolving an information need differently, either through the type of search engine they use, the queries they enter, the documents they view and so on. By bringing people together to search at the same time, and enabling a SoK across collaborators, we can enable group members to benefit from the diversity across the group.

An interesting aspect of IR is that it is often a learning process [38]. Users can approach an IR engine with a vague notion of their information need and develop this notion through interactions with a search engine. Therefore in SCIR an SoK policy needs to enable users to benefit from the discoveries of their collaborators during the search, in order to ensure that each group member can develop their perceptions of the shared task.

8

DoL and SoK policies can be achieved either through a user-driven co-ordination of the search, where users themselves divide the task and share information, or from a system driven coordination of the search task, where the back-end system attempts to coordinate the task across multiple users. To date the vast majority of work in SCIR has focussed on designing for awareness across collaborating searchers in the hope that the users themselves can coordinate the search. In the next section we will review some of this work.

4. Designing for Awareness in SCIR

SCIR "in the wild" has been observed in co-located settings, where groups of people either arrange themselves around a single computer terminal with one user acting as the driver and other users contributing, or in remote settings in which users work independently on their own browsers and communicate via instant messaging or emailing.

When collaborating over a search task in a co-located setting, users are often working in close proximity, awareness of the activities of others is naturally high but sometimes at the expense of group productivity as users can be distracted by the actions of others. In a remote setting, users can search more independently through their own browser, however in such a setting awareness amongst group members often comes at a high-cost, such as having to email relevant links to colleagues [39].

SCIR systems to-date have attempted to improve group productivity in both settings through the provision of awareness cues. The majority of work to-date has been in supporting remote SCIR.

4.1. Remote SCIR

GroupWeb [40] is an early example of a remote SCIR system. In Group-web the web browser was used as a "group presentation tool" in which one master browser selected a webpage and this page was displayed across all collaborators' browsers. Users could attach annotations to webpages and these annotations were shared across all collaborators. GroupWeb represented a tightly coupled collaborative browsing experience, awareness was high across collaborators, however such tight coupling did not allow for a DoL across collaborators as all users are required to view the same information at the

same time. The system did provide for a SoK across users via shared annotations and users could use this tool to discuss the contents of a document and share their opinions of it with others.

Building on from Groupweb, several other remote SCIR systems were developed to enable a more relaxed form of synchronous search [41, 42]. The W4 browser [40] contained many of the awareness widgets seen in state-of-the-art SCIR systems including:

- Session history – which showed the URLs of documents seen by group members

- Shared bookmarks – which showed the URLs of documents marked as relevant to the search by group members

- Chat window – for communicating ideas

- Shared whiteboard – for group brainstorming

Users could also embed chat-sessions, links and annotations directly on any webpage. The relaxed searching experience and awareness rich environment of the W4 browser can enable collaborators to coordinate their search task while searching independently. The group could use W4 to achieve a DoL by assigning certain tasks to each collaborator using the chat mechanism. Furthermore, recording those documents seen by group members in the session history widget can allow users to skip over those documents already viewed by their collaborators. SoK can be achieved through brainstorming ideas across the shared whiteboard, or viewing the annotations made by others on webpages. Users could also benefit from the discoveries of their collaborators through the shared bookmarked area where they could see those documents deemed relevant to the search task. This bookmark widget coupled with the session history is a good example of how DoL and SoK policies can work together in tandem: users can choose to ignore viewing those documents contained in the history conscious that such documents that were seen by others *and* deemed relevant to the search will be saved in the bookmarked area.

SearchTogether [43, 44] is an Internet Explorer plugin that allows for remote SCIR across groups of searchers. The system was developed to explicitly support awareness and division of labour across collaborators. *Query awareness* is achieved by showing each group member's name and photo

alongside their query history. Page-specific awareness allows users to see visitation counts, ratings and comments associated with web pages. On a ranked list of search results, users can see those webpages that have been viewed by their collaborators. Both the query history and the page specific awareness widgets can achieve a DoL and SoK across searchers. As noted by the authors, users can avoid unnecessary duplication of effort by being aware of those query terms entered by collaborators, furthermore users can learn from seeing those query terms entered by others. Showing visitation information directly in the search results pane, rather than in a separate history widget, requires less effort on the part of collaborators to avoid unnecessary duplication of effort. A similar approach known as *group member URL traversal awareness* (GMUTA) was proposed by [45]. SearchTogether also enables users to recommend webpages to their collaborators. This facility is similar in purpose to the bookmarked area of systems like W4, but through a recommendation device items of relevance are "pushed" towards users rather than being "pulled" by the users themselves. Furthermore, the recommendation mechanism enables documents to be pushed to particular users rather than to the group as a whole, such a mechanism could support specialisation in the group search where each user is assigned particular aspects of the search task to investigate.

4.2. Co-Located SCIR

Recently the development of both ubiquitous computing devices and shared displays has facilitated the development of co-located collaborative information retrieval systems [46, 47, 48, 49, 50]. By bringing users together in a face-to-face environment, these systems improve the awareness across collaborating searchers inherently. A co-located collaborative environment is an awareness-rich setting, collaborators can converse easily while they search, furthermore this communication is enhanced by non-verbal communication such as the facial-expressions, postures and gestures of their search partners. If utilised effectively, this increased awareness has the potential to enable a more effective division of labour and a greater sharing of knowledge across the collaborating group. Design decisions through the choice of interfaces and widgets on these systems can further improve awareness across users, such as including sounds associated with actions or exaggerated gestures to complete different tasks [29].

Han et. al [47] proposed a system for co-located collaborative browsing on mobile devices. The system divided a webpage into distinct regions and

distributed these sections across user's mobile devices. Maekwa et. al [48] extended this approach to allow for a more intelligent division of the webpage, by considering factors such as the device the users was working on (e.g. mobile phone, laptop) and the user's interests represented in a set of keywords, when dividing the page.

In [49] we presented Físchlár-DiamondTouch , a system for co-located collaborative search on a DiamondTouch multi-user touch-sensitive electronic tabletop [51]. The system enabled two users to sit opposite each other and search for video shots by querying the corpus with either text or content-based visual queries and interacting with the search results using point and drag touch-gestures. Two variants of the system were developed, one where the focus of the design was on supporting individual efficiency through the search by providing rapid menu-items and another which attempted to increase the awareness across collaborating users. Awareness was supported through the provision of a single set of search widgets for users (i.e. a single text query box and a single area to move saved shots to etc.) this required users to drag shots to different parts of the table, thereby making their collaborators aware of their activities in their peripheral vision. Each significant action also had an associated sound cue: for example saving a shot caused a "stamping" noise to be heard along with a visual "saved" stamp appearing on the shot. Results from our experiments showed that users both preferred and performed better using the awareness based system.

Físchlár-DiamondTouch supported division of search results implicitly: results from queries were orientated towards the nearest users, users were therefore inclined to interact with those shots closest to them. Due to the provision of only one query box, users collaborated on formulating a shared query with each user suggesting terms to enter. Users also collaborated over relevance judging whereby if one user was unsure if a particular video shot was relevant they consulted their search partner.

Collaboration in Físchlár-DiamondTouch was very much front-loaded, although users could co-ordinate the DoL and SoK of the search to a certain extent, the group search was impeded by the fact that the system was using a standard, single-user, IR backend. For example, the provision of only one query box inhibited a parallel search, rather users were tightly coupled around a single shared group query. We felt that such a scenario could be improved through the development of IR algorithms and methods that were designed specifically to support multi-user search. In the next section we will discuss our work in developing system-mediated DoL and SoK techniques.

5. System-Mediated Techniques for SCIR

In our work we have explored the potential for system-mediated DoL and SoK in SCIR search [15, 52, 53, 30]. Our motivation for this is based on the belief that an SCIR system can be improved with a back-end search engine which is specifically designed to support collaboration. This belief is based on both our own experiences with Físchlár-DiamondTouch and from our review of the state-of-the-art awareness based systems for SCIR. In essence, in order to achieve an effective group search using a purely awareness-based system, collaborators need to both search *and* coordinate a shared group task. For example, in order to achieve an effective division of labour and ensure duplication of effort is removed, collaborating users would need to coordinate their actions using an instant messaging client for example or check if the link they are about to follow has already been followed by their collaborators. For sharing of knowledge, users need to read those documents saved in the bookmarked folder, or recommended to them. As observed by [23] requiring users to move their attention between the search task and the coordination of the group effort may require too much of a user's cognitive load or at best be an unwelcome distraction from the search task.

Through implementing system-mediated techniques, the burden for co-ordination is moved from the users to the back-end, and this allows users to concentrate on the search task itself without concern for organising the group effort. To this end we developed system-mediated techniques for both DoL and SoK and in this section we will summarise our work to-date.

5.1. System-Mediated Division of Labour

In [30] we presented some light-weight back-end division of labour techniques for adhoc SCIR. These techniques were based upon two assumptions. Firstly, that users searching together to satisfy the same information need will enter the same or similar queries and as a result, have returned to them ranked lists which contain many of the same documents, and this has been shown to be the case [30]. Secondly that we could remove those documents for all subsequent ranked lists given that at least one member of the group had already seen it. Considering that all users are searching to satisfy the same information need, this approach seems reasonable.

In our experiments with simulated search we investigated the effects of different levels of division of labour on an SCIR session. In particular, we examined the effects of a full DoL environment in which no two users were

13

presented with the same documents. This was implemented by removing from a user's ranked lists those documents that had been viewed by others previously, and also those that were contained on another collaborator's ranked list that they were examining. Our experimental results showed that an SCIR system with a system-mediated DoL policy was better than an SCIR system without any such coordination. The experiments also highlighted the importance of a DoL for SCIR, as those SCIR system variants without any DoL were shown to perform significantly worse than two independent search baselines where users were simulated searching without any collaboration.

These results were very encouraging, considering the simplicity of the methods involved and showed the potential and importance for system-mediated DoL. The advantage of removing redundant documents, rather than just highlighting as is the case in an awareness-based SCIR system, is that redundant documents in the ranked list are replaced with new unseen documents. This enables the group to cover a greater amount of the collection over the course of the search. A potential downside of such an approach is the possibility than in removing relevant documents from a user's ranked list we could impede a user's natural development of their understanding of the information need. One way to overcome such an issue could be the inclusion of a separate awareness-based bookmark window so that they could review these relevant documents that were found by others. Another issue with a system-mediated DoL technique could occur when a user who is not familiar with the search topic skips over a potentially relevant document, causing it to be removed from the subsequent ranked lists of the group. One solution here could be to down-weight the scores of documents that have been seen by others rather than using the brute-force approach of removing them altogether.

5.2. Sharing of Knowledge

In [52, 30] we presented our techniques for system-mediated sharing of knowledge for SCIR. Our approach was to develop methods that enabled users to benefit from the discoveries of their collaborators during the search and in this way attempt to improve the performance of the group activity. This work was motivated by the observation that while many state-of-the-art SCIR systems encouraged users to make explicit judgments in the form of bookmarks, no attempt was made to use these relevance judgements in the search process itself. Relevance feedback (RF) is an IR method proven to improve the quality of ranked lists by reformulating a user's query automatically

in the light of relevance judgments [54]. This is achieved through expanding the query with new terms from these relevance judgments (query expansion) and reweighting query terms to take into account their distribution in these relevant documents (term reweighting).

If we have an environment where two or more users are providing relevance judgments to an SCIR system, then, when performing relevance feedback for a user, the SCIR system has an opportunity to incorporate each user's relevance judgments into a the relevance feedback process, which has the potential to improve the quality of the ranked lists returned to each user.

In our work we investigated two strategies for sharing relevance information across collaborating searchers: *collaborative* and *complementary* relevance feedback. By sharing judgments in this way, users can benefit from the efforts of their collaborators in finding these relevant documents.

5.2.1. Collaborative Relevance Feedback

In [52] we outlined our approach for combining relevance information in a collaborative relevance feedback method.

In particular we extended the probabilistic relevance feedback method [55], in order to allow for a weighted combination of each collaborator's relevance judgments (bookmarks) in the process. In this way the reformulated query contains relevance information from multiple collaborating users thus allowing for an implicit sharing of knowledge across users.

In our experiments we investigated if sharing relevance judgments across users in an SCIR environment improved the performance of the group search. We performed simulated experiments using the same procedure as used for our DoL experiments. Our results showed that a collaborative relevance feedback method provided small improvements in the group's performance over the entire search, with greater improvements achieved at the early part of the search.

Synchronous collaborative information retrieval systems, by their very nature, bring together multiple collaborating users, each with a certain level of expertise and experience. Some users may be more familiar with a topic than others and this may be reflected in the quality of their relevance judgments during the group search activity For example, a novice user may not understand the search topic entirely and therefore may be mistaken in their relevance assessments. As our proposed collaborative relevance feedback algorithm operates by combining the relevance judgments of all collaborators poor relevance assessments unless recognised and dealt with may degrade the

15

quality of results being presented to users.

In [30] we investigated the effects of poor relevance assessments on the performance of a collaborative relevance feedback process. Unsurprisingly our results showed that poor relevance judgments decreased the performance of the collaborative relevance feedback process significantly. Such a result represents a significant concern for any process which attempts to combine relevance information from many users collaborating together.

In [30] we proposed an authority weighting mechanism which used the capacity for weighted combination in the collaborative relevance feedback formula to bias the process in favour of the more authoritative user. Our experimental results showed that by biasing the collaborative relevance feedback process towards more authoritative users the performance of the collaborative relevance feedback mechanism is improved significantly.

5.2.2. Complementary Relevance Feedback

Our experiments in [52] showed that a collaborative relevance feedback method provided only modest increases in performance over an entire SCIR search. Our analysis of the data revealed that the collaborative relevance feedback process was causing the ranked lists returned to each collaborating user to become very similar. This result is intuitive given that the process allowed relevance information from one user to affect the process for their search partners. We felt that this loss of diversity was likely to be a key factor in the poor performance of our collaborative relevance feedback method.

In [30] we investigated another use of relevance judgments, in particular through a complementary relevance feedback method. The idea behind complementary relevance feedback is to perform relevance feedback in such a way that diversity is maintained or increased across users' ranked lists. We experimented with two methods here: complementary query expansion and clustering of the relevance judgment set. Complementary query expansion ensured that when performing relevance feedback, each user's reformulated relevance feedback query contained unique terms. The idea here that these distinct queries should produce ranked lists with increased diversity.

Our clustering approach was to partition the set of relevance judgments into k partitions where k was the number of users collaborating in the search, the idea being that each cluster would contain a set of relevance judgments that were as diverse as possible from each other.

Although both complementary query expansion and clustering improved the diversity across user's ranked lists, neither improved the performance of

an SCIR search. This result was disappointing, however we believe there is still potential for a complementary feedback method.

5.3. iBingo: Mobile Collaborative Video Search

In [53] we implemented our system-mediated concepts for division of labour and sharing of knowledge in "iBingo" a collaborative video search system developed for Apple iPhones. Through iBingo we explored the concepts of DoL and SoK through *information trails*. Using simple touch gestures users could perform rapid content-based searching, which allowed users to build an information trail through the collection. The back-end search engine implemented a DoL by ensuring that once a video shot had been examined by one user it was removed from the trails of their collaborators. SoK was achieved through the sharing of user-trails and was performed when a user had exhausted their own information trail. In this way, when both users were pursuing their own trails a DoL ensured no duplication of effort, whilst a SoK policy ensured that users could help each other when needed. iBingo also represented our first attempt at combining system-mediated and awareness coordination, the users were working in close proximity and were free to discuss the search topic while the back-end coordinated their actions.

6. Discussion

In this paper we discussed the application of two CSCW concepts: division of labour and sharing of knowledge, to a synchronous collaborative information retrieval environment. We presented two approaches by which these concepts can be realised in SCIR: either through the provision of awareness cues or through system-mediated techniques. Both approaches have advantages and disadvantages associated with them.

6.1. Awareness versus System-Mediated Coordination

For certain search topics, a clear division of the task across group members may not be possible using user-user awareness coordination. If users are searching in the same information space then they will invariably be returned documents that have already been viewed by their collaborators. Our system-mediated DoL techniques would remove such documents from a user's ranked list thereby ensuring that no two users spend time viewing the same documents and ensure a complete DoL. The downside of removing relevant documents from a user's ranked list before returning it to them is that we

17

may disrupt the development of a user's understanding of the information need. Users may think that the queries they have entered are poor because they are seeing no relevant documents returned, when in fact the query could have returned many documents but they were removed because they had already been seen by a collaborator. In such a scenario an awareness widget could inform a user that a certain number of documents have been removed, with the option to show these if needed.

The bookmarks and recommendation widgets of awareness-based SCIR systems allow for a sharing of relevant documents across collaborating users. If users wish to view these documents they can follow the links and view these documents. In our system-mediated collaborative and complementary relevance feedback approaches we attempt to enable users to benefit from the discoveries of their collaborators without having to view the documents themselves and mark them as relevant. The benefit here is in the form of improved ranked lists, however what these system-mediated techniques fail to capture is the benefit, in terms of their understanding of the information need, a user gains from seeing the relevant document. A combination of both approaches would allow users to benefit from the relevance judgments of their collaborators whilst having the option to view these relevant documents if they wish. Furthermore, upon reviewing a document marked by another collaborator, a user may elect to agree or disagree with the user's relevance judgment, and this could reinforce a collaborative or complementary relevance feedback algorithm.

To summarise, although users can attempt to coordinate the search tasks amongst themselves through an awareness-based system, such techniques could impede the group search performance, while a pure system-mediated approach may cause users to feel lost in the search. Our conclusion is that the inherent benefits and shortcomings of each would suggest that a state-of-the-art SCIR system should endeavour to support both approaches.

6.2. Future work in System-Mediated Techniques

In our work we have explored the notion of authority in relation to the relevance feedback process. For future work it would be interesting to extend this notion of authority, and model cases where users can skip relevant material. Modelling such an environment would enable us to investigate how a division of labour policy would operate in an environment where searchers may read a document but fail to recognise it as relevant. At present, the DoL techniques would exclude this document from all ranked lists returned to all

users for the rest of the search. This would cause the group to miss relevant material due to the actions of a poor searcher. A potential solution to this issue would be to extend the division of labour policy to allow for an approach which reduced the rank of seen documents rather than excluding them, as is currently the case. This dampening effect could be weighted by the perceived expertise of a user so that, if a more expert user reads a document without providing a relevance judgment on it, then we can be more confident that the document is non-relevant than if a poor searcher performed the action.

We proposed techniques which allowed for a user-biased weighting of relevance judgments, and we investigated one application of a user-biasing through authority weighting. However we believe there are many more applications of a user-biased collaborative relevance feedback process. For example, in a real search system, a user could use a user-biased relevance feedback process to favour their own relevance judgements over their search partners', thereby allowing the results to be tailored to them. Or a user may decide to use an inverted approach and bias their results in the favour of their search partner if they feel they cannot locate any relevance documents on their own. We believe there are many more applications of our proposed collaborative relevance feedback process, which we have not even considered.

We introduced the notion of imperfect relevance judgements to synchronous collaborative search. In traditional, single-user information retrieval, the notion of imperfect relevance judgments is less of an issue as, if a user has made a relevance judgment, then it should be considered as relevant for that user. When we move to a synchronous collaborative domain, in which a group of users are searching together, the issue of non-relevance or misunderstanding of the search topic could have a major effect of the performance of an SCIR search. We believe that this also represents an interesting area from future research. We investigated the application of authority weighting, and as discussed in the previous section we believe there is scope for future research here and, in particular, in future work we would like to apply some of the various query performance predictions techniques to the area of SCIR authority weighting. Although the results of our evaluations of our proposed complementary relevance feedback techniques were not favourable, we believe there is still scope for further research here. In particular we feel that the clustering of documents and terms in a SCIR search may prove useful. In this thesis we attempted to cluster documents and terms into two clusters, one for each user. An interesting avenue for further research would be to try to use clustering to discover unique concepts in the search topic. For ex-

ample a search topic Wildlife Extinction may have concepts related to zoos, poachers, animals, etc. By recognising the underlying concepts present in the relevance judgments, we may be able to make a more intelligent division of the clusters across users, whereby each user is assigned a unique concept to search.

Division of labour and sharing of knowledge are two closely related concepts in SCIR, we therefore believe that both should be realised in all systems. For example, a division of labour policy can remove documents it deems to be redundant to the search given that they have been seen by other users, as a sharing of knowledge policy ensures that such documents that are also marked as relevant, through bookmarks, are brought to the attention of the group.

7. Conclusion

In this paper we introduced and discussed two concepts which we believe are vital to the development of effective SCIR systems and on which we carried out experimental investigations: division of labour and sharing of knowledge. We outlined how both techniques had been recognised as important in the CSCW domain and discussed how they applied to the emerging research field of SCIR. We discussed how both can be achieved either through the provision of awareness or by our system-mediated techniques and showed how they can improve the performance of SCIR. Our suggestion is that future SCIR systems should endeavour to support both techniques through the provision of awareness cues and system-mediated techniques. Through iBingo we have implemented such a system, and in future work we will pursue further integration of awareness and system-mediated DoL and SoK techniques.

References

[1] Clarence A. Ellis, Simon J. Gibbs, and Gail Rein. Groupware: some issues and experiences. *Commun. ACM*, 34(1):39–58, 1991.

[2] Saul Greenberg and Mark Roseman. GroupWeb: a WWW browser as real time groupware. In *CHI '96: Conference companion on Human factors in computing systems*, pages 271–272, Vancouver, British Columbia, Canada, 1996. ACM Press.

[3] Barry Smyth, Evelyn Balfe, Jill Freyne, Peter Briggs, Maurice Coyle, and Oisín Boydell. Exploiting Query Repetition and Regularity in an Adaptive Community-Based Web Search Engine. *User Modeling and User-Adapted Interaction*, 14(5):383–423, 2005.

[4] J.-W. Ahn, P. Brusilovsky, and R. Farzan. Investigating users' needs and behaviour for social search. In *Proceedings of the Workshop on New Technologies for Personalized Information Access (part of the 10th International Conference on User Modelling (UM '05))*, pages 1–12, Edinburgh, Scotland, UK, July 2005.

[5] Del.icio.us. http://del.icio.us/, last acceessed May 2009.

[6] Heystaks. http://www.heystaks.com/, last accessed May 2009.

[7] Robert W. Root. Design of a multi-media vehicle for social browsing. In *CSCW '88: Proceedings of the 1988 ACM conference on Computer-supported cooperative work*, pages 25–38, Portland, Oregon, United States, 1988. ACM Press.

[8] T. D. Wilson. On user studies and information needs. *Journal of Documentation*, 37(1):3–15, 1981.

[9] Nicholas C. Romano Jr., Dmitri Roussinov, Jay F. Nunamaker, and Hsinchun Chen. Collaborative information retrieval environment: Integration of information retrieval with group support systems. In *HICSS '99: Proceedings of the Thirty-Second Annual Hawaii International Conference on System Sciences-Volume 1*, Maui, Hawaii, 1999. IEEE Computer Society.

[10] Michael B. Twidale, David M. Nichols, and Chris D. Paice. Browsing is a collaborative process. *Inf. Process. Manage.*, 33(6):761–783, 1997.

[11] Steven Poltrock, Jonathan Grudin, Susan Dumais, Raya Fidel, Harry Bruce, and Annelise Mark Pejtersen. Information seeking and sharing in design teams. In *GROUP '03: Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work*, pages 239–247, Sanibel Island, Florida, USA, 2003. ACM.

[12] Preben Hansen and Kalervo Järvelin. Collaborative information retrieval in an information-intensive domain. *Inf. Process. Manage.*, 41(5):1101–1119, 2005.

[13] Jonathan Foster. Collaborative information seeking and retrieval. *Annual Review of Information Science and Technology*, 40(1):329–356, 2006.

[14] Meredith Ringel Morris. A survey of collaborative web search practices. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1657–1660, New York, NY, USA, 2008. ACM.

[15] Colum Foley, Alan F. Smeaton, and Hyowon Lee. Synchronous collaborative information retrieval with relevance feedback. In *Collaborate-Com 2006 - 2nd International Conference on Collaborative Computing: Networking, Applications and Worksharing, pp1-4*, Adelaide, Australia., 2006.

[16] Kari Kuutti. The concept of activity as a basic unit of analysis for cscw research. In *ECSCW'91: Proceedings of the second conference on European Conference on Computer-Supported Cooperative Work*, pages 249–264, Norwell, MA, USA, 1991. Kluwer Academic Publishers.

[17] Ke-Thia Yao, Robert Neches, In-Young Ko, Ragy Eleish, and Sameer Abhinkar. Synchronous and asynchronous collaborative information space analysis tools. In *ICPP '99: Proceedings of the 1999 International Workshops on Parallel Processing*, page 74, Washington, DC, USA, 1999. IEEE Computer Society.

[18] David Benyon, Phil Turner, and Susan Turner. *Designing Interactive Systems - People, Activities, Contexts, Technologies.* Addison-Wesley, 2004.

[19] C. Gutwin, S. Greenberg, R. Blum, J. Dyck, K. Tee, and G. McEwan. Supporting informal collaboration in shared-workspace groupware. 14(9):1411–1434, 2008.

[20] Paul Dourish and Sara Bly. Portholes: supporting awareness in a distributed work group. In *CHI '92: Proceedings of the SIGCHI conference*

on Human factors in computing systems, pages 541–547, Monterey, California, United States, 1992. ACM.

[21] Andrew Cockburn, Andrew Cockburn, Saul Greenberg, and Saul Greenberg. Making contact: Getting the group communicating with groupware. pages 31–41. ACM Press, 1993.

[22] Roy Want, Andy Hopper, Veronica Falc ao, and Jonathan Gibbons. The active badge location system. *ACM Trans. Inf. Syst.*, 10(1):91–102, 1992.

[23] John Adcock, Jeremy Pickens, Matthew Cooper, Lisa Anthony, Francine Chen, and Pernilla Qvarfordt. FXPAL Interactive Search Experiments for TRECVID 2007. In *TRECVid2007 - Text REtrieval Conference TRECVID Workshop*, Gaithersburg, MD, USA., November 2007.

[24] M.R. Morris. Collaborating alone and together: Investigating persistent and multi-user web search activities. Technical report, Microsoft Research, January 2007.

[25] S Talja. Information sharing in academic communities: types and levels of collaboration in information seeking and use. *The New Review of Information Behaviour Research*, 3:143–159, 2002.

[26] M Sharples. Adding a little structure to collaborative writing. *CSCW in Practise:An Introduction and Case Studies*, 1993.

[27] Thomas W. Malone and Kevin Crowston. What is coordination theory and how can it help design cooperative work systems? pages 357–370, 1990.

[28] Steven Poltrock and Mark Handel. Modeling collaborative behavior: Foundations for collaboration technologies. *Hawaii International Conference on System Sciences*, 0:1–10, 2009.

[29] Alan F. Smeaton, Colum Foley, Cathal Gurrin, Hyowon Lee, and Sinéad Mc Givney. Collaborative Searching for Video Using the Físchlár System and a DiamondTouch Table. In *TableTop2006 - The 1st IEEE International Workshop on Horizontal Interactive Human-Computer Systems, pp149-156*, pages 149–156, Adelaide, Australia, 2006.

[30] Colum Foley and Alan F. Smeaton. Synchronous collaborative information retrieval: Techniques and evaluation. In *ECIR 2009 - 31st European Conference on Information Retrieval*, pages 42–53, Toulouse, France, 2009.

[31] Steve Whittaker, Erik Geelhoed, and Elizabeth Robinson. Shared workspaces: how do they work and when are they useful? *Int. J. Man-Mach. Stud.*, 39(5):813–842, 1993.

[32] John C. Tang. Findings from observational studies of collaborative work. *Int. J. Man-Mach. Stud.*, 34(2):143–160, 1991.

[33] Kum-Yew Lai, Thomas W. Malone, and Keh-Chiang Yu. Object lens: a "spreadsheet" for cooperative work. *ACM Trans. Inf. Syst.*, 6(4):332–353, 1988.

[34] Victoria Bellotti and Sara Bly. Walking away from the desktop computer: distributed collaboration and mobility in a product design team. In *CSCW '96: Proceedings of the 1996 ACM conference on Computer supported cooperative work*, pages 209–218, New York, NY, USA, 1996. ACM.

[35] Keith J Dilley. Informatics & analytics - the team approach. IRFS2008, the Information Retrieval Foundation Symposium, November 2008.

[36] Michael J. Muller, Werner Geyer, Beth Brownholtz, Eric Wilcox, and David R. Millen. One-hundred days in an activity-centric collaboration environment based on shared objects. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 375–382, New York, NY, USA, 2004. ACM.

[37] Randall H. Trigg, Lucy A. Suchman, and Frank G. Halasz. Supporting collaboration in notecards. In *CSCW '86: Proceedings of the 1986 ACM conference on Computer-supported cooperative work*, pages 153–162, New York, NY, USA, 1986. ACM.

[38] Peter Ingwersen. *Information Retrieval Interaction*. Taylor Graham, London, 1992.

[39] M.R. Morris and S Amershi. Shared sensemaking: Enhancing the value of collaborative web search tools. In *CHI 2008 Workshop on Sensemaking*, Florence, Italy, 2008.

[40] S. Gianoutsos and J. Grundy. Collaborative work with the World Wide Web: adding CSCW support to a Web browser. In *Procedingss of Oz-CSCW'96, DSTC Technical Workshop Series, University of Queensland*, University of Queensland, Brisbane, Australia, August 1996.

[41] G. Cabri, L. Leonardi, and F. Zambonelli. Supporting Cooperative WWW Browsing: a Proxy-based Approach. In *7th Euromicro Workshop on Parallel and Distributed Processing*, pages 138–145, University of Maderia, Funchal, Portugal, Feb 1999. IEEE Press.

[42] R. Krishnappa. Multi-user search engine (MUSE): Supporting collaborative information seeking and retrieval. Master's thesis, University of Missouri-Rolla, University of Missouri-Rolla, Rolla, 2005.

[43] Meredith Ringel Morris and Eric Horvitz. SearchTogether: an interface for collaborative web search. In *UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology*, pages 3–12, Newport, Rhode Island, USA, 2007. ACM Press.

[44] SearchTogether. http://research.microsoft.com/en-us/um/redmond/projects/searchtogether/, last accessed May 2009.

[45] Efstratios T. Diamadis and George C. Polyzos. Efficient cooperative searching on the web: system design and evaluation. *Int. J. Hum.-Comput. Stud.*, 61(5):699–724, 2004.

[46] Alan F. Blackwell, Mark Stringer, Eleanor F. Toye, and Jennifer A. Rode. Tangible interface for collaborative information retrieval. In *CHI '04: extended abstracts on Human factors in computing systems*, pages 1473–1476, Vienna, Austria, 2004. ACM.

[47] Richard Han, Veronique Perret, and Mahmoud Naghshineh. WebSplitter: a unified XML framework for multi-device collaborative Web browsing. In *CSCW '00: Proceedings of the 2000 ACM conference on Computer supported cooperative work*, pages 221–230, Philadelphia, Pennsylvania, United States, 2000. ACM.

[48] Takuya Maekawa, Takahiro Hara, and Shojiro Nishio. A collaborative web browsing system for multiple mobile users. In *PERCOM '06: Proceedings of the Fourth Annual IEEE International Conference on Pervasive Computing and Communications (PERCOM'06)*, pages 22–35, Washington, DC, USA, 2006. IEEE Computer Society.

[49] Alan F. Smeaton, Hyowon Lee, Colum Foley, and Sinéad Mc Givney. Collaborative Video Searching on a Tabletop. *Multimedia Systems Journal*, 12(4):375–391, 2006.

[50] Meredith Ringel Morris, Andreas Paepcke, and Terry Winograd. TeamSearch: Comparing Techniques for Co-Present Collaborative Search of Digital Media. In *TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, pages 97–104, Washington, DC, USA, 2006. IEEE Computer Society.

[51] Paul Dietz and Darren Leigh. DiamondTouch: a multi-user touch technology. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 219–226, Orlando, Florida, 2001. ACM.

[52] Colum Foley, Alan F. Smeaton, and Gareth J. F. Jones. Combining relevance information in a synchronous collaborative informationretrieval environment. *Collaborative and Social Information Retrieval and Access Techniques for Improved User Modelling*, 2008.

[53] Alan F. Smeaton, Colum Foley, Daragh Byrne, and Gareth Jones. ibingo mobile collaborative search. In *CIVR 2008 - ACM International Conference on Image and Video Retrieval. VideOlympics @ CIVR*, 2008.

[54] Donna Harman. Relevance feedback revisited. In *SIGIR '92: Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 1–10, Copenhagen, Denmark, 1992. ACM Press.

[55] S. E. Robertson and K. Spärck Jones. Relevance weighting of search terms. *Journal of the American Society for Information Science*, 27(3):129–146, 1976.