



HAL
open science

User-Driven System-Mediated Collaborative Information Retrieval

Laure Soulier, Chirag Shah, Lynda Tamine

► **To cite this version:**

Laure Soulier, Chirag Shah, Lynda Tamine. User-Driven System-Mediated Collaborative Information Retrieval. Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval, Jul 2014, Gold Coast, Australia. 10.1145/2600428.2609598 . hal-01132594v2

HAL Id: hal-01132594

<https://hal.science/hal-01132594v2>

Submitted on 29 Jan 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

User-Driven System-Mediated Collaborative Information Retrieval

Laure Soulier
Universite de Toulouse UPS
IRIT
118 route de Narbonne
F-31062 Toulouse, France
soulier@irit.fr

Chirag Shah
School of Communication &
Information (SC&I)
Rutgers University
4 Huntington St, New
Brunswick, NJ 08901, USA
chirags@rutgers.edu

Lynda Tamine
Universite de Toulouse UPS
IRIT
118 route de Narbonne
F-31062 Toulouse, France
tamine@irit.fr

ABSTRACT

Most of the previous approaches surrounding collaborative information retrieval (CIR) provide either a user-based mediation, in which the system only supports users' collaborative activities, or a system-based mediation, in which the system plays an active part in balancing user roles, re-ranking results, and distributing them to optimize overall retrieval performance. In this paper, we propose to combine both of these approaches by a role mining methodology that learns from users' actions about the retrieval strategy they adapt. This hybrid method aims at showing how users are different and how to use these differences for suggesting roles. The core of the method is expressed as an algorithm that (1) monitors users' actions in a CIR setting; (2) discovers differences among the collaborators along certain dimensions; and (3) suggests appropriate roles to make the most out of individual skills and optimize IR performance. Our approach is empirically evaluated and relies on two different laboratory studies involving 70 pairs of users. Our experiments show promising results that highlight how role mining could optimize the collaboration within a search session. The contributions of this work include a new algorithm for mining user roles in collaborative IR, an evaluation methodology, and a new approach to improve IR performance with the operationalization of user-driven system-mediated collaboration.

Categories and Subject Descriptors

H.3.3 [INFORMATION STORAGE AND RETRIEVAL]: Information Search and Retrieval - *Clustering, Retrieval models, Search process*

General Terms

Collaborative Information Retrieval

Keywords

Collaborative information retrieval, role mining, user study

1. INTRODUCTION

In an information retrieval (IR) process, there are several ways to improve retrieval performance or effectiveness. Some approaches, such as pseudo-relevance feedback [13] and query expansion [30] focus on improving IR systems and algorithms. Others try to improve users' experience by offering related queries [2] or making personalized recommendations [29]. In the end, the effectiveness of an IR setting has to do with user's preferences and/or skills just as much as it has with system performance. One way that effectiveness could be improved that is often overlooked is through collaboration. The approaches described above are based on an individual searching, retrieving, and evaluating information. But what if multiple people get involved in these processes? Beyond simply dividing up the work, there could be other potential benefits of multiple people working together on a search-focused project. Specifically, if these people bring in a variety of skills, then we have an opportunity to achieve synergic effects [22]. The question is how do we know who has which skills and how do we leverage them. We could let the users decide it and let the system be simply a supporting structure. This is often referred to as user-mediated collaborative information retrieval (CIR) [15, 23]. The other option is to have the system explicitly assign prior roles and optimizing retrieval-related processes accordingly. This is often referred to as a system or an algorithmically-mediated CIR [20, 24].

Both of these lines of work have their merits and limitations. In this paper, we propose an approach that attempts to combine the best of both of these techniques. Specifically, instead of imposing the roles onto the users (system-mediated) or having the users explicitly come up with their roles (user-mediated), we propose to mine roles of collaborators implied within a collaborative search session. We refer to this as user-driven system-mediated CIR. More particularly, the contributions of this paper are twofold:

- A *role mining methodology*, which learns how the participants in a collaboration are different and uses that to suggest roles.
- An *experimental evaluation* based on two user studies involving a total number of 70 collaborative search sessions, which shows that role mining optimizes the collaboration within the search process.

In what follows, we introduce, in Section 2, related work surrounding the CIR domain, and more particularly, the consideration of user roles within this context. Section 3 focuses on the research questions and our role mining methodology. Section 4 details our experimental setup based on two interactive IR studies and Section 5 presents and discusses the obtained results. Finally, Section 6 concludes the paper and presents future work.

2. RELATED WORK

2.1 User-based mediation and system-based mediation for CIR

A collaborative search setting involves multiple users aiming at collaboratively solving a shared information need [4, 6, 25]. In [25] authors showed how collaborative search could act as a front-end for existing search engines and re-rank results based on the learned preferences of a community of users. They attempted to demonstrate this concept by implementing the I-Spy system [4]. However, this is often ineffective or non-trivial. For instance, Hyldegard [8], with her studies of information seeking and retrieval in a group-based education setting, found that even though people in a collaborative group to some extent demonstrated similar cognitive experiences as the individuals in Kuhlthau’s Information Search Process (ISP) model [11], these experiences did not only result from information seeking activities, but also from work-task activities and intra-group interactions. Her further work also indicated [9] that group based problem solving is a dynamic process that shifts between a group perspective and an individual perspective.

Complexity in CIR appears mainly from the fact that collaboration is characterized by four main dimensions [6]: (1) the *intent*, depending on if collaboration is implicit or explicit, (2) the *level of mediation*, relying on how much the system is active, (3) the *concurrency*, analyzing if the search session is synchronous or not, (4) the *location*, describing a session involving remote or co-located users. Rising challenges in CIR [10] underline the practical necessity of building interfaces adapted to collaboration and from an IR perspective, leveraging IR techniques and models for supporting collaboration. To tackle these challenges, two main lines of work can be distinguished: user-based mediation and system-based mediation.

Previous user-based mediation work have proposed collaborative interfaces [16, 23] supporting collaboration within a search process by using devices, such as interactive tables [16], or tools, such as shared workspaces [23]. The second category of work, referred to a system-oriented mediation, relies on the incorporation of symmetric user’s roles [3, 17] or asymmetric ones within the IR process [20, 24, 26] with the aim of enhancing the collaboration and by the way, the overall IR performance.

In [6], authors proposed a role taxonomy where several sets of roles have been introduced: (1) symmetric (peers) [6] in which users are considered as having symmetric roles; (2) domain A expert and domain B expert [6, 26] in which users are characterized as experts of different domains; (3) search expert and search novice, respectively domain expert and domain novice [6, 27] in which users have different famil-

ilarity levels towards a domain, or search tools; search expert and domain expert [6] in which one user has access to more functions and features with search tools, whereas the other one is more familiar with the domain knowledge; (4) Prospector and Miner [6, 20] in which Prospector’s task is devoted to formulating search requests, whereas the Miner’s one consists of selecting relevant documents; Gatherer and Surveyor [24], in which the Gatherer aims at quickly scanning documents for identifying relevant ones whereas the Surveyor’s task is devoted to favoring diversity.

2.2 Role mining

Several research studies [21, 1, 5, 28] highlighted the value of roles as means of control in collaborative applications. Role is viewed as a tool for better understanding users’ actions and interactions by putting them in context. For instance, Rodden [21] identified the value of control in computer-supported cooperative work systems and listed a number of underlying projects with their corresponding schemes for implementing control. Golder and Donath [5] considered Usenet environment for analyzing newsgroup communities. They proposed a role taxonomy involving two roles which they posit, namely “Celebrities” and “Ranters”, and three well-know roles, namely “Lurkers”, “Trolls” and “Newbies”. They also highlighted the underlying strategies of each role and relationships between roles. Role-based control gives rise to a challenging research question: how can we mine dynamically latent roles from user-group actions and interactions in a collaborative setting?

Two main approaches for role mining are proposed in the literature. The first one, rising from the social network perspective, relies on the analysis of the topology properties of users’ network through information interaction [18, 12, 19, 7]. Nowicki and Snijders [18] presented a block-structure mining method based on pair-wise graph relations using Gibbs sampling. In [12], authors identified microblog networks leaders by applying PageRank algorithm on the followers network while Pal et al. [19] proposed to apply a probabilistic clustering method that uses a set of features from microblogging practices. Henderson et al. [7] provided an unsupervised learning approach for role mining relying on the transformation of feature-based multi-dimensional matrix where nodes represent users and edges represent information stream. The underlying objective is to select the better compression model that summarizes the users’ behavior leading to roles. Different from the first approach, the second one discovers user’s roles by looking at the similarities vs. dissimilarities between interaction data distributions among users [14]. Authors proposed, first, the LDA-based Author-Recipient-Topic (ART) model which is a global role mining approach applied in a communication social network for highlighting emerging roles through users’ behavior similarities. Second, they proposed a pair-wise focused version of the ART by introducing the Role-Author-Recipient-Topic (RART) model which emphasizes roles of the author and the recipient according to the message-contents, assuming that these roles could be fixed or could evolve throughout the message writing process.

Based on relevant literature, we can see that different forms of roles have been introduced within a collaborative context, leveraging users’ skills and experiences in order to en-

hance the collaboration. However, while some research studies focused on mining latent roles from users’ interactions, previously proposed system-mediated collaborative ranking models assign explicitly predefined roles to users regardless of the skills in which they are the most effective. In the work reported here, we attempt to combine insights from the above works to bridge between system-mediated and user-mediated collaboration and show how to leverage specific latent roles of collaborators from users’ actions while relaxing the assumption behind pre-assigned roles. To the best of our knowledge, this is the first attempt for mining users’ roles within a collaborative search, which enables implicitly and dynamically assigning roles to users in which they can be most effective at the current search stage.

3. RESEARCH DESIGN

3.1 Research Questions

In this work presented here, we do not assign any roles to users in the beginning of the session. Users’ roles are emphasized by users’ actions and are mined throughout the synchronous collaborative search session. Keeping this in mind, we address in this work the following research questions:

- RQ1. If the searchers working in collaboration are different in some aspects, how are they different?
- RQ2. How do we infer users’ roles by taking into account their actions within a collaborative search session?
- RQ3. How can we use role information to improve CIR and IR processes in general?

3.2 Mining Users’ Roles

We initially consider a user-driven search setting in which a set of users are working to meet a shared information need by identifying relevant documents.

The goal of our approach is to analyze how users are different, and, accordingly suggest roles given their differences in order to optimize their collaborative outcome. In this way, the search session can be viewed as a sequence of two steps: (1) a user-driven level in which we assume that people start out with the same roles and responsibilities, but then over time, they may exhibit differences and affine their search behavior; and (2) a system-mediated level in which predefined roles have been identified, and the system optimizes the collaborative search process. In what follows, we first define preliminary notions, and, then describe our role mining methodology. Note that in order to simplify our explanation and economize our experiments, we will consider only two participants in a collaborative search process.

3.2.1 Preliminaries

Let us consider a typical scenario of collaborative IR where a pair of users (u_1, u_2) collaborate during a time interval T in order to answer a shared information need I . Each user acts separately through the synchronous collaborative search session S by formulating his own queries belonging to the shared information need I . For clarity, we introduce the following notions and notations:

User’s search session. The user search session S_u^t of user u consists of all previous queries $q_u^1, q_u^2, \dots, q_u^{t-1}$ ordered by time, up to timestamp t ($0 < t < T$), and their corresponding documents results sets $D_u^1, D_u^2, \dots, D_u^{t-1}$. Since the most important factor affecting the role mining task is naturally how much the users are different while searching for information during the CIR process, we favor a feature-based representation of the user’s search session. Our goal is to characterize the user through his past implicit actions’ features $F = \{f_1, \dots, f_k, \dots, f_n\}$ (eg. number of visited pages, dwell time, number of formulated queries etc.) captured during his/her search process. To achieve this, S_u^t is represented as a matrix $S_u^{(t)} \in \mathcal{R}^{t \times n}$ where each element $S_u^{(t)}(t_l, f_k)$ represents the cumulative value of feature f_k for user u aggregated across the search session during the time interval $[t_0 \dots t_l]$. This aggregation highlights the overall behavior of the user at timestamp t_l since looking at feature values at a particular timestamp may not be enough considering that individual behaviors are constantly changing. Moreover, it avoids the bias induced by noisy search actions.

Roles. A pair $R_{1,2}$ of roles (r_1, r_2) involves a couple of users (u_1, u_2) assumed to have different specific search features.

We introduce a role pattern $P^{R_{1,2}} \in P$, where P is the role pattern pool, as a set of two elements:

- A feature-based matrix $F^{R_{1,2}} \in \mathcal{R}^{F \times F}$, which represents the underlined behaviors of users of the collaborative group if they are performing these roles. Each element $F^{R_{1,2}}(f_j, f_k)$ highlights the expected relation such as positive correlation (+1), negative correlation (-1) or independence (0) between the users’ behavior features f_j, f_k belonging to the pair of roles under consideration. Here, we assume that a pair of roles may be mined from a collaborative search session only if users’ behaviors are different and opposite. For instance, the Gatherer and Surveyor roles [24] are differentiated considering the fact that the Gatherer seeks highly relevant documents and thus formulates queries with a higher overlap than the Surveyor whereas the Surveyor spends less time on Webpages favoring the exploration of the information space, and accordingly the average value of query success is lower, than the Gatherer. Simply having behavior differences for these two features is not enough because what if the same person is better at both? Therefore, a negative correlation between these behaviors is considered to ensure that both the behaviors are found differently in both the users. Accordingly, we can define the related pair of roles through: (1) negative correlation between the query word overlap and spent time on Webpages features, (2) negative correlation between query word overlap and query success features. Thus, we define the pattern kernel \mathcal{K} as the set of feature pairs highlighted as negatively correlated where $\mathcal{K}^{R_{1,2}} = \{(f_j, f_k); F^{R_{1,2}}(f_j, f_k) = -1\}$.
- A role attribution function $Role(S_{u_1}^{(t)}, S_{u_2}^{(t)}, R_{1,2})$, which assigns each role of the pair $R_{1,2}$ to users u_1 and u_2 , according to his/her intrinsic search behaviors $S_{u_1}^{(t)}$ and

$S_{u_2}^{(t)}$ respectively. For this purpose, we compare values $S_{u_1}^{(t)}(t_l, f_k)$ and $S_{u_2}^{(t)}(t_l, f_k)$ of at least one search feature f_k of both collaborators u_1 and u_2 for identifying which role of the pair $R_{1,2}$ is the most likely suitable for each user. For instance, in the case of Gatherer-Surveyor roles, the query overlap feature should be higher for the Gatherer. The output of this document attribution function is a tuple of ordered roles (r, r') where the first element is the role assigned to user u_1 , and the second element is the role assigned to user u_2 . We define the role attribution function $Role(S_{u_1}^{(t)}, S_{u_2}^{(t)}, R_{1,2})$ as next.

$$\begin{aligned} \text{Role} : (S_{u_1}^{(t)}, S_{u_2}^{(t)}, R_{1,2}) &\rightarrow \{(r_1, r_2), (r_2, r_1)\} \\ (S_{u_1}^{(t)}, S_{u_2}^{(t)}, R_{1,2}) &\mapsto \begin{cases} (r_1, r_2) & \text{if } S_{u_1}^{(t)}(t_l, f_k) > S_{u_2}^{(t)}(t_l, f_k) \\ (r_2, r_1) & \text{otherwise} \end{cases} \end{aligned} \quad (1)$$

Finally, our objective is to mine at each timestamp t the role pattern $P^{R_{1,2}}$ that likely maps the users' past actions represented through the users' search sessions $S_{u_1}^{(t)}$ and $S_{u_2}^{(t)}$, and accordingly to assign detected roles to collaborators.

3.2.2 Methodology

We now introduce our methodology to mine the pair of users' roles involved in a synchronous collaborative search session. At each regular timestamp t_l , a role mining method, illustrated in Algorithm 1, is launched and runs in two main stages, as described in what follows.

Algorithm 1: Role mining in collaborative search setting

Data: $S_{u_1}^{(t)}, S_{u_2}^{(t)}, P$

Result: $\{(r_1, r_2), (r_2, r_1)\} \in R$

begin

```

/* Step 1: Identifying users' search behavior
differences */
for each feature  $f_k \in F$  do
   $p_{u_1, u_2}^{(t_l)}(f_k) = \text{ks.test}(S_{u_1}^{(t)}(f_k), S_{u_2}^{(t)}(f_k));$ 
/* Step 2: Identifying users' roles */
for each feature  $f_j \in F$  do
  for each feature  $f_k \in F$  do
    if  $(p_{t_l}^{u_1, u_2}(f_j) < \theta)$  and  $(p_{t_l}^{u_1, u_2}(f_k) < \theta)$  then
       $C_{u_1, u_2}^{(t_l)}(f_j, f_k) = \rho(\Delta_{u_1, u_2}^{(t_l)}(f_j), \Delta_{u_1, u_2}^{(t_l)}(f_k));$ 
    else
       $C_{u_1, u_2}^{(t_l)}(f_j, f_k) = 0;$ 
 $R_{1,2}^* = \text{argmin}_{R_{1,2}} \|F^{R_{1,2}} \ominus C_{u_1, u_2}^{(t_l)}\|$ 
subject to:
 $\forall_{(f_j, f_k) \in \mathcal{K}^{R_{1,2}}} F^{R_{1,2}}(f_j, f_k) - C_{u_1, u_2}^{(t_l)}(f_j, f_k) > -1;$ 
if  $R_{1,2}^* \neq \text{null}$  then
  return  $Role(S_{u_1}^{(t)}, S_{u_2}^{(t)}, R_{1,2}^*);$ 
return null;

```

1. *Identifying users' search behavior differences.* Keeping in mind that a pair of users' roles implies different users'

search actions, the first stage consists of capturing significant differences between the users' search actions through the analysis of their search sessions. More formally, in order to identify the level of difference between both users u_1 and u_2 with respect to each feature $f_k \in F$, we use the non-parametric Kolmogorov-Smirnov statistical test, noted $\text{ks.test}()$ in Algorithm 1, between vectors $S_{u_1}^{(t_l)}(f_k)$ and $S_{u_2}^{(t_l)}(f_k)$. The obtained p -value for the statistical test is noted as $p_{u_1, u_2}^{(t_l)}(f_k)$. Significant differences imply that users u_1 and u_2 have different search skills or strategies and thus could play different roles in the search session.

2. *Identifying users' roles.* Our objective in this stage is to characterize the difference between the users' search strategies in contrast to the previous stage that only identifies their difference. For this aim we attempt to detect candidate role patterns through user's search actions. The method we adopted consists of computing the correlation between the user's search features. More specifically, we consider features by pairs and analyze the correlation of users' differences towards the considered pairs of features.

Computing the correlation between the feature differences values rather than the feature values themselves avoids prior assignment of any role to each user. Thus, we compute the correlation between the vectors $\Delta_{u_1, u_2}^{(t_l)}(f_j)$ and $\Delta_{u_1, u_2}^{(t_l)}(f_k)$ with $(f_j, f_k) \in F \times F$. We build a difference matrix $\Delta_{u_1, u_2}^{(t_l)} \in \mathcal{R}^{t_l \times F}$ of each feature f_k where each element $\Delta_{u_1, u_2}^{(t_l)}(t_l, f_k)$ is defined as the difference of the value of user u_1 and user u_2 towards the search feature f_k at timestamp t_l :

$$\Delta_{u_1, u_2}^{(t_l)}(t_l, f_k) = S_{u_1}(t_l)(t_l, f_k) - S_{u_2}^{(t)}(t_l, f_k) \quad (2)$$

These correlation computations rest on the condition that the significance of both implied search feature differences between both users is highlighted by the Kolmogorov-Smirnov statistical test; otherwise we set the correlation value to 0. Finally, a correlation matrix $C_{u_1, u_2}^{(t_l)} \in \mathcal{R}^{F \times F}$ is build as:

$$C_{u_1, u_2}^{(t_l)}(f_j, f_k) = \begin{cases} \rho(\Delta_{u_1, u_2}^{(t_l)}(f_j), \Delta_{u_1, u_2}^{(t_l)}(f_k)) & \text{if } ((p_{t_l}^{u_1, u_2}(f_j) < \theta) \\ & \text{and } (p_{t_l}^{u_1, u_2}(f_k) < \theta)) \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where θ is a significance level.

Afterwards, the role pattern is detected in the case that:

- The involved search features highlighted by a negative correlation within the role pattern are significantly different and negatively correlated between both users. Indeed, we assume that users might behave significantly differently at least for features involved within the role pattern. This condition can be formulated by the fact that if a feature correlation measure $F^{R_{1,2}}(f_j, f_k)$ connected to the role pattern P^{r_1, r_2} equals to -1 , then the correlation $C_{u_1, u_2}^{(t_l)}(f_j, f_k)$ underlying the users' search behaviors between features f_j and f_k should strictly be lower than 0.

Table 1: User study statistics.

| | US1 | US2 |
|--|----------------|----------------|
| Topic | Gulf oil spill | Global warming |
| Number of pairs | 60 | 10 |
| Total number of visited documents | 4326 | 1935 |
| Total number of saved/bookmarked documents | 254 | - |
| Total number of saved snippets | 485 | 208 |
| Total number of submitted queries | 921 | 313 |
| Average number of words by query | 5.45 | 4.73 |
| Average repeatability of query words | 182.04 | 30.28 |

- The role pattern correlation matrix is the most likely similar to the collaborative group correlation matrix. We are aware that an exact matching between correlation matrices respectively relying on role pattern and users’ search behaviors is difficult to reach since the role pattern is characterized by negative correlations equal to -1. Therefore, a reasonable role-based identification is to assign the role pattern correlation matrix $F^{R_{1,2}}$ which is the most similar to the one $C_{u_1, u_2}^{(t_i)}$ obtained for the pair of users (u_1, u_2) at timestamp t_i .

More formally, a pair $R_{1,2}$ of roles (r_1, r_2) is assigned to a pair of user (u_1, u_2) at timestamp t_i if it satisfies the following conditions:

$$\underset{R_{1,2}}{\operatorname{argmin}} \|F^{R_{1,2}} \ominus C_{u_1, u_2}^{(t_i)}\| \quad (4)$$

subject to :

$$\forall_{(f_j, f_k) \in \mathcal{K}^{R_{1,2}}} F^{R_{1,2}}(f_j, f_k) - C_{u_1, u_2}^{(t_i)}(f_j, f_k) > -1$$

where $\|\cdot\|$ represents the Frobenius norm and \ominus is the minus operator defined as:

$$F^{R_{1,2}}(f_j, f_k) \ominus C_{u_1, u_2}^{(t_i)}(f_j, f_k) = \begin{cases} F^{R_{1,2}}(f_j, f_k) - C_{u_1, u_2}^{(t_i)}(f_j, f_k) \\ \quad \text{if } F^{R_{1,2}}(f_j, f_k) \in \{-1; 1\} \\ 0 \quad \text{otherwise} \end{cases}$$

Once the pair $R_{1,2}$ of roles has been detected for a collaborative group involving users u_1 and u_2 , the respective roles are assigned to collaborators according to the role attribution function $\operatorname{Role}(S_{u_1}^{(t)}, S_{u_2}^{(t)}, R_{1,2})$.

4. EXPERIMENTAL SETUP

Experimental evaluation is based on two collaborative IR user studies. This section describes these studies, the CIR tasks, the data collected, and the evaluation metrics used.

4.1 Data Acquisition and Processing

4.1.1 User Studies

We consider two user studies, involving pairs of users, recruited at Rutgers University. All of them are students and got compensation for their participation. The search tasks consisted of writing a report on an exploratory topic. Throughout the whole session, participants had the possibility to refine their information need by exploring and reformulating queries dealing with the topic. Each task was supported by a collaborative information system [23], which includes a chat system and search tools enabling participants to bookmark Webpages and save snippets. Participants were

allowed to compile snippets for building the final report. In the background, the system recorded search log data, such as visited and bookmarked pages, saved through snippets and submitted queries. Statistics about the both datasets are presented in Table 1. We notice that the repeatability of query words represents the average occurrence number of each word included in a query over the whole set of queries. For instance, in the dataset US1, a query includes in average 5 words and each of them are generally used within 182 other queries. This underlines the fact that queries are mainly issued from a reformulation process. Below, we highlight the specificities of each user study:

US1. This study involved 60 dyads working on the exploratory topic of “Gulf oil spill”. Participants were given 25 minutes for synchronously working on the exploratory search task and writing the report. The participants were randomly assigned to co-located or remote settings.

US2. This study involved 10 dyads in which participants are distinct to those implied in the first user study. The study setup was similar to the previous one with the following exceptions: (1) the exploratory task focused on “global warming”; (2) the collaborative information system allowed saving of snippets and not pages as bookmarks; (3) participants disposed of 30 minutes for solving the exploratory task; and (4) participants disposed of co-located computers.

4.1.2 Document Dataset

In order to perform the system-mediated setting, we built a unique document index for the two user studies including: (1) visited pages of the whole set of participants of both user studies; and (2) top 100 Google search results of queries submitted by the whole set of participants. The top search results are extracted using the Google Custom Search API¹. We, thus, obtained a dataset with 74,844 distinct Webpages. The datasets were then parsed to extract title and main content from the Webpages.

4.1.3 Search Session Features

From search log data, we extracted three categories of features, summarized in Table 2, characterizing participants’ search actions. The first category analyzes the query submission process of each participant, for instance by emphasizing if he is better devoted to formulating queries depending on his intent level of submitting a query, modeled by his query number or if he is looking for diversity with respect to his used vocabulary within the successive submitted queries,

¹<https://developers.google.com/custom-search/>

Table 2: Features used to represent a user search session.

| Category | Description | Measurement |
|------------------------|-----------------------------|--|
| Query-based features | Number of queries | Number of submitted queries |
| | Query length | Average number of tokens within queries |
| | Query success | Average ratio of successful pages (time spent over 30s) over queries |
| | Queries overlap | Average ratio of shared words number among successive queries |
| Page-based features | Number of pages | Number of visited pages |
| | Number of pages by query | Average number of visited pages by query |
| | Page dwell time | Average time spent between two visited pages |
| Snippet-based features | Number of snippets | Number of snippets |
| | Number of snippets by query | Average number of snippets by query |

modeled by his query overlap value. In addition, we also identified features based on query length and query success. We assume that a query is successful if the retrieved pages are visited for over 30 seconds. The second category analyses browsing behavior related to Webpages. For instance, we can analyze with the page dwell time if one participant spends a lot of time for reading one page or not. The number of visited pages, whether normalized by the number of submitted queries or not, shows if a participant is more willing to read a lot of documents. The last category analyses the snippet-based browsing behaviors emphasizing if users are more willing to find interesting documents. During our role mining methodology, we estimate the feature value for a $[0; t]$ time-window for each timestamp t .

4.2 Roles

In accordance to the literature review, we consider the two role-based collaborative ranking models modeling the Gatherer/Surveyor roles [24] and the Prospector/Miner roles [20].

The Gatherer and Surveyor roles are different considering the fact that the Gatherer seeks for highly relevant documents and thus formulates queries with a higher overlap than the Surveyor whereas the Surveyor spends less time on Webpages favoring the exploration of the information space, and accordingly the average value of query success is lower, than the Gatherer. Thus, we can define the related role pattern through: (1) negative correlation between the query word overlap and spent time on Webpages features; (2) negative correlation between query word overlap and query success features. The role attribution function underlying this couple of roles specifies that the user with the higher query overlap feature value will be assigned to the role of Gatherer whereas the other one will be detected as the Surveyor.

The main difference between the roles of Prospector and Miner remains on the task performed by each user. The Prospector’s task is more devoted for reformulating queries in order to get result diversity whereas the Miner is looking for relevant documents. Consequently, we can expect for these roles to have a negative correlation between query overlap and number of submitted queries features. The underlying role attribution function specifies that the user characterized by the most important number of submitted queries will be assigned to the role of Prospector whereas the other one will be detected as the Miner.

4.3 Protocol Design and Baselines

We divided the collaborative search sessions into regular timestamps of sixty seconds enabling us to detect users’ search actions differences, and accordingly to predict their roles within collaboration. We perform the following evaluation protocol:

1. For each timestamp t , mine roles that users are performing using our methodology detailed in Section 3.
2. If a role pattern is detected, use the queries submitted between timestamps t and $t + 1$ with these different models:
 - **BM25-CIR**: the BM25 baseline ranking model on each single query submitted by one of the two users. This scenario aims at simulating the user-driven session in which users perform collaboration by means of independent search engines.
 - **GS-CIR**: The Gatherer/Surveyor baseline collaborative ranking model in order to evaluate the impact of considering a system-mediated collaborative ranking model, namely the Gatherer/Surveyor one [24] not necessarily connected with users’ search behaviors.
 - **PM-CIR**: The Prospector/Miner baseline collaborative ranking model in order to evaluate the impact of considering a system mediated collaborative ranking model, namely the Prospector/Miner one [20] not necessarily connected with users’ search behaviors.
 - **Ra-CIR**: The random baseline mining model in order to evaluate the impact of a random version of a user driven system mediated CIR setting in which roles are randomly mined or not. For this purpose, we randomly launch either the BM25 model, the Prospector/Miner one [20] or the Gatherer/Surveyor one [24].
 - **RB-CIR**: Role-based model as proposed in this paper. The model requires a couple of queries, respectively for each user. Thus, at each time a query is submitted by a user, we identify the last query submitted by the other user. Both queries are inputs of the CIR model. In the case of any role is detected, we launched the BM25 ranking model.

4.4 Metrics

Table 3: Users’ search behavior average differences. **absDiff**: absolute difference. **p-value** ($p < 0.05$): number of collaborative groups with significantly different members: *: at least one group. **: at least 50% of collaborative groups. **: at least 75% of collaborative groups. ****: all the groups.

| | | Timestamp | 1 | 2 | 3 | 4 | 5 | 10 | 15 | 20 | 25 | 30 |
|-----|---------|-----------|--------|--------|--------|--------|--------|--------|--------|--------|--------|-------|
| US1 | QS | absDiff | 0.535 | 0.746 | 0.801 | 0.837 | 0.820 | 1.004 | 1.023 | 1.101 | 1.240 | - |
| | | p-value | | ** | ** | *** | *** | *** | **** | **** | **** | - |
| | QN | absDiff | 1.153 | 1.339 | 1.746 | 1.932 | 2.339 | 2.898 | 3.508 | 3.983 | 4.190 | - |
| | | p-value | | * | ** | ** | *** | *** | *** | *** | *** | - |
| | QO | absDiff | 0.309 | 0.340 | 0.413 | 0.421 | 0.380 | 0.291 | 0.303 | 0.259 | 0.238 | - |
| | | p-value | | * | * | ** | *** | *** | **** | **** | **** | - |
| DWP | absDiff | 52.725 | 67.169 | 76.289 | 87.344 | 82.109 | 80.544 | 72.196 | 56.760 | 52.567 | - | |
| | p-value | | ** | *** | *** | *** | *** | **** | **** | **** | **** | - |
| US2 | QS | absDiff | 0.333 | 0.467 | 0.548 | 0.699 | 0.661 | 0.811 | 0.986 | 0.570 | 0.400 | 0.107 |
| | | p-value | | ** | ** | *** | *** | *** | *** | *** | *** | *** |
| | QN | absDiff | 1.200 | 2.200 | 2.100 | 2.500 | 2.700 | 3.100 | 5.800 | 6.500 | 7.000 | 0.500 |
| | | p-value | | * | ** | ** | ** | *** | *** | **** | **** | *** |
| | QO | absDiff | 0.276 | 0.380 | 0.355 | 0.316 | 0.387 | 0.231 | 0.224 | 0.157 | 0.134 | 0.144 |
| | | p-value | | * | ** | ** | ** | **** | **** | **** | **** | **** |
| DWP | absDiff | 36.558 | 34.285 | 28.116 | 23.635 | 17.386 | 13.435 | 13.988 | 13.062 | 12.554 | 12.779 | |
| | p-value | | *** | *** | *** | *** | *** | *** | *** | *** | *** | |

For each dataset, we build the ground truth by considering the whole set of distinct documents which have been visited by at least two different users as relevant. Using this process, we obtained 159 and 47 relevant documents for dataset US1 and US2 respectively.

We used collaborative-based metrics aiming at emphasizing the retrieval-based synergic effect of the collaboration [22]. Considering our CIR setting, we computed these metrics at the timestamp level as follows:

- The precision $Prec@R(g)$ at rank R of a collaborative group g [22]:

$$\begin{aligned}
 Prec@R(g) &= \frac{1}{T(g)} \sum_{t=1}^{|T(g)|} Prec@R(g)^{(t)} \\
 &= \frac{1}{T(g)} \sum_{t=1}^{|T(g)|} \frac{RelCov@R(g)^{(t)}}{Cov@R(g)^{(t)}} \quad (5)
 \end{aligned}$$

where $T(g)$ is the whole set of the timestamps for which document lists have been retrieved. Coverage $Cov@R(g)^{(t)}$, respectively the relevant coverage $RelCov@R(g)^{(t)}$, of a collaborative group g is the total number of distinct top R retrieved documents, respectively distinct top R relevant retrieved documents, at a timestamp t for the whole set of participants.

- The recall $Recall@R(g)$ at rank R of group g [22]:

$$\begin{aligned}
 Recall@R(g) &= \frac{1}{T(g)} \sum_{t=1}^{|T(g)|} Recall@R(g)^{(t)} \\
 &= \frac{1}{T(g)} \sum_{t=1}^{|T(g)|} \frac{RelCov@R(g)^{(t)}}{|RelDoc|} \quad (6)
 \end{aligned}$$

where $|RelDoc|$ expresses the total number of relevant documents for the exploratory task.

- The F-measure $F_{syn}@R(g)$ at rank R of a collaborative group g [22]:

$$F@R(g) = \frac{1}{T(g)} \sum_{t=1}^{|T(g)|} \frac{2 * Prec@R(g)^{(t)} * Recall@R(g)^{(t)}}{Prec@R(g)^{(t)} + Recall@R(g)^{(t)}} \quad (7)$$

5. RESULTS

In this section we present and discuss results, divided in three categories: those measuring differences between users, those analyzing the extracted roles, and those evaluating the effectiveness of role mining for CIR.

5.1 Analyzing Users’ Differences

Regarding the first stage of our methodology, we analyzed the evolving user search actions using the features presented in Table 2. Table 3 illustrates the average absolute differences observed between collaborators of each group according to each user study with respect to features involved in role patterns presented in Section 4.2, namely the query success (QS), the number of submitted queries (QN), the query overlap (QO) and the dwell time of visited pages (DWP) at each timestamp, from 1 to 30.

We can see from this table that, for the whole set of features, absolute differences among collaborators of at least one group are highlighted as positive and significant in the beginning of the session after the first minute. Therefore, significance of collaborators’ search behavior differences arises the assumption that users are different early in the session and provide the foundation to our role mining methodology. This statement is reinforced by the fact that from 75% to 100% of the collaborative groups are characterized by significant differences between collaborators after 5 minutes.

5.2 Analyzing the Mined Role Changes

In this section, we study how our proposed role mining methodology acts in the collaborative search setting. For this purpose, we define three metrics aiming at measuring

the changes of mined roles within a search session:

1. The **number of role couples**, noted $N_{couples}$, normalized by the number of timestamps, respectively 25 and 30 for US1 and US2, for analyzing the stability of users' behavior differences throughout the search session. More particularly, it measures the number of changes within chronologically mined couples of roles, for which we add the couple of role mined for the first timestamp. We assume that there is a change when, from timestamp t to timestamp $t + 1$, the roles assigned to the pair of users are different. A value close to 1 means that the detected couples of roles are different for each two successive timestamps while a value close to 0 reflects that the pair of users have constant behaviors throughout the search process.

2. The **number of distinct roles**, noted N_{roles} , assigned to a particular user, throughout the search session also normalized by the number of timestamps. We consider five possibilities of individual roles: Prospector, Miner, Gatherer, Surveyor or nothing. We note that the maximum value, highlighting that a user is detected performing each of these five roles in a session, is 0.2 for US1 and 0.16 for US2.

3. The **couple-role ratio**, noted R_{change} , which helps to analyze if, for the same detected couple of roles at two different timestamps within the search session, the same role is assigned to the same user or if behaviors have changed to what extend roles of collaborators became opposite. For instance, a user might be Prospector in the beginning of the session and Miner at the end. For this purpose, we compute the ratio R_{change} of the number of role couples $N_{couples}$ and the number of distinct roles N_{roles} for each collaborative search session. A value equals to 1 highlights that, even if the assigned couple of roles to the pair of users varies and is redundant throughout the session, roles within a couple are steadily assigned to the same user. In other words, when a user is detected as a Gatherer at one timestamp t of the session, if the Gatherer-Surveyor couple is again mined at timestamp t' , this user will be still detected as a Gatherer. A value lower than 1 for a collaborative group reflects that, for at least one couple of roles detected at two different timestamps, each user received successively the different associated roles, e.g. Gatherer at timestamp t and Surveyor at timestamp t' . In contrast, a ratio value is lower than 1 means that the couples of roles detected for the group varies more than the evolution of users' search behaviors when the same couple of roles is detected.

Table 4 gives statistics of these metrics for both datasets. More particularly, we computed the average value and the range, under the format "average [$min - max$]".

Table 4: **Analysis of the evolution of mined roles throughout the search session.**

| | US1 | US2 |
|---------------|--------------------|-------------------|
| $N_{couples}$ | 0.13 [0.04 – 0.28] | 0.08[0.03 – 0.2] |
| N_{roles} | 0.09[0.04 – 0.16] | 0.09[0.04 – 0.16] |
| R_{change} | 0.88[0.29 – 1.5] | 1.18[0.5 – 1.5] |

We can see from Table 4 that, for both user studies, 1) the low number of changes in terms of couples $N_{couples}$ and distinct roles N_{roles} and 2) the couple-role ratio value R_{change} very close to 1, highlight a certain stability in roles assigned, first, to the pair of users, and second, to the individual users themselves.

Indeed, the normalized number $N_{couples}$ of couples, respectively 0.13 and 0.08 for US1 and US2, reflects that general search behaviors of the pairs are stable over the session. We notice that participants' behaviors of the first user study US1 are less stable than the second one US2 as we detected more changes in terms of mined couples of roles. Both user studies have equivalent results in terms of distinct role number mined for individual user with a value equal to 0.09. Thus, statements are slightly opposite, although both close to 1, with respect to the couple-role ratio for both user studies: pairs of participants in US1 are slightly more variant in terms of global search behaviors than individual ones which means to what extend we detected more mined couples of roles than for the roles assignment for a solely participant. In contrast, in US2, users' intrinsic search behaviors fluctuate more than mined couples of roles in so far as that, for the same couple of roles detected at two different timestamps, the same user was labelled differently.

This analysis lets us explore if the time window of the role detection impacts the retrieval effectiveness of such CIR scenario. To this aim, we estimated the retrieval effectiveness of our role-based CIR scenarios, noted $RB-CIR$ by computing the role mining methodology for a time window from 1 to 5 minutes. Table 5 illustrates the obtained results considering the F-measure. Beforehand, we highlight that metric values are low due to the small number of relevant documents building the relevance judgement for each dataset. We can notice from Table 5 that, for the user study US2, even if the retrieval effectiveness seems to be higher for a 5-minutes time window, results are not significantly different due to the low number of groups. However, in the first user study, namely US1, the optimal time window is 1 minute. The analysis of student statistical tests highlights that the 1-minute time window-based retrieval effectiveness is significantly higher ($p < 0.05$) than the other settings; and settings based on 2, 3, 4 or 5 minutes are not statistically significant between each other. This shows that, even if the mined roles role-couples do not vary significantly a lot, it remains however more effective to propose a mining for each minute in order to better fit users' search behaviors and therefore optimize the collaboration within the search session. For this reason, in the remainder of the paper, we consider a 1-minute time window for mining roles of collaborators.

Table 5: **Impact of the role mining time window on the retrieval effectiveness.**

| | 1 | 2 | 3 | 4 | 5 |
|-----|----------------|---------|---------|---------|----------------|
| US1 | 0.01673 | 0.01650 | 0.01654 | 0.01649 | 0.01647 |
| US2 | 0.07061 | 0.06683 | 0.06931 | 0.06650 | 0.07203 |

We also performed a correlation analysis among the couple-role ratio and the F-measure. Results showed non significant correlations among the involved variables: respectively 0.003 and -0.003 for both dataset US1 and US2. This highlights

Table 6: Comparative the impact of role mining on the retrieval effectiveness of a CIR task. %Chg; *RB-CIR* improvement. Student test significance *: $0.01 < t \leq 0.05$; **: $0.001 < t \leq 0.01$; ***: $t \leq 0.001$.

| | | <i>Prec@20</i> | | | <i>Recall@20</i> | | | <i>F@20</i> | | |
|-----|----------|----------------|---------|-----------------|------------------|---------|-----------------|-------------|---------|-----------------|
| | | value | %Chg | <i>p</i> -value | value | %Chg | <i>p</i> -value | value | %Chg | <i>p</i> -value |
| US1 | BM25-CIR | 0.041 | 10.408 | * | 0.010 | 4.636 | * | 0.016 | 5.372 | |
| | GS-CIR | 0.038 | 18.316 | *** | 0.008 | 25.504 | *** | 0.014 | 24.521 | *** |
| | PM-CIR | 0.050 | -9.482 | | 0.012 | -13.991 | | 0.019 | -13.397 | |
| | Ra-CIR | 0.041 | 11.484 | * | 0.009 | 12.895 | * | 0.015 | 12.777 | * |
| | RB-CIR | 0.045 | - | | 0.010 | - | | 0.017 | - | |
| US2 | BM25-CIR | 0.075 | 3.347 | | 0.063 | 2.586 | | 0.069 | 2.833 | |
| | GS-CIR | 0.058 | 34.636 | | 0.040 | 63.818 | * | 0.046 | 52.786 | * |
| | PM-CIR | 0.092 | -16.051 | | 0.078 | -16.493 | | 0.084 | -16.317 | |
| | Ra-CIR | 0.070 | 10.714 | | 0.056 | 16.201 | | 0.062 | 14.324 | |
| | RB-CIR | 0.077 | - | | 0.065 | - | | 0.071 | - | |

that the retrieval effectiveness of our proposed CIR setting based on role mining does not depend on the number of distinct mined roles for a user, but rather on adapting the ranking model to users’ search behaviors.

5.3 Effectiveness Evaluation

5.3.1 Evaluating the Global Retrieval Effectiveness

Here, we measure and discuss the impact of role mining on the overall retrieval effectiveness of a collaborative search session. To this aim, we average, for each metric, their values over the whole set of participant groups. Significance level used in our methodology is fixed to $\alpha = 0.05$.

Table 6 shows the comparative results between the retrieval effectiveness of our role-based CIR setting, noted *RB-CIR*, and the four baseline settings, namely *BM25-CIR*, *GS-CIR*, *PM-CIR* and *Ra-CIR*, described in Section 4.3.

Results highlight significant improvements for most of the collaborative-based retrieval measures. For user study US1, mining users’ roles significantly enhance the retrieval effectiveness of the overall session from +10.408% to +18.316% for the precision metrics, except a non significant decrease for the Prospector-Miner-based CIR scenario. In contrast, the improvements for user study US2 are less significant due to the fact that this study involves less collaborative sessions and therefore, the significance of student test is more difficult to reach. Nevertheless, we obtained significant improvements from 52.786% to 63.818% with respect to the *GS-CIR* baseline setting.

More generally, we notice the following general trends for both datasets:

- Leveraging from role mining within a CIR search session overpasses the aggregation of an individual scenario highlighted by the *BM25-CIR* baseline. This shows the necessity of leveraging from collaborative search settings by analyzing users’ differences and mining their roles in order to optimize the collaboration.
- Our role-based CIR scenario provides better results than the scenario relying only on the Gatherer-Surveyor collaborative ranking model. This highlights that our proposed scenario is better adapted to the constant setting of the Gatherer-Surveyor CIR model

which only merges and splits search results retrieved from the last submitted queries of users.

- The Prospector-Miner algorithm seems to provide better results than our scenario, however these improvements are not significant for both datasets. One possible explanation of an equivalent retrieval effectiveness as our proposed scenario might be that the Prospector-Miner CIR model analyzes the retrieved documents and submitted queries from the beginning of the session for providing optimized document rankings. Similarly, our role mining methodology also rests on users’ actions from the beginning of the session for proposing a better adapted CIR ranking model.
- Our role mining methodology seems to be more effective than a scenario in which roles are assigned randomly. This emphasizes the reliability of our methodology for mining roles which seems to accurately match users’ search behaviors, and accordingly optimize the collaboration.

Thus, we can conclude that mining roles within a collaborative search session optimizes the collaboration and improves the synergic effect of collaborators.

5.3.2 Evaluating the Timestamp-based Retrieval Effectiveness

Our objective here is to analyze more in-depth retrieval effectiveness across time of our proposed methodology-based scenario, noted *RB-CIR*, and the four baseline settings, namely *BM25-CIR*, *GS-CIR*, *PM-CIR* and *Ra-CIR*. Thus, for each timestamp t , we average over all participant groups their F-measure at rank R for timestamp t , estimated by combining $Prec@R(g)^t$ and $Recall@R(g)^t$, as detailed in Equation 7. Figure 1 illustrates the retrieval effectiveness for the five settings according to each timestamp of the search session. For the first user study, namely US1, the curve representing system-driven approach surrounded by role mining, namely *RB-CIR*, generally overpasses the baseline curves. This statement is more slightly identifiable in Figure 1b due to the low number of collaborative groups detected as performing a couple of roles for a particular timestamp. Therefore, differences between *RB-CIR* and baseline settings are more variable and non-significant. Nevertheless, we can conclude that mining roles generally optimizes the collaboration




Figure 1: Comparative effectiveness throughout the collaborative search session.

throughout the search process. Moreover, graphics highlight that the effectiveness of our method seems to decrease at the end of the session. It can be explained by the fact that sessions were fixed and quite long, between 25 and 30 minutes, and we noticed during user studies that participants were less involved in the task at the end of the session. We can assume, therefore, that the cognitive effort for reformulating queries was higher which may negatively impact the information need expression, and therefore, the reliability of mined roles in this unstable context.

6. CONCLUSION AND FUTURE WORK

Collaboration can be a useful strategy in many situations, including information retrieval. As researchers have shown [20], using collaboration one could achieve something greater than the sum of individual searchers' contributions. A way to achieve this synergy is by leveraging diverse sets of skills brought by different people [22].

In this paper, we proposed a user-driven system-mediated approach for mining organically emerging user roles and using that information to improve upon various aspects of retrieval performance including relevance and novelty. Specifically, we presented a new role mining methodology that helps us understand how two searchers collaborating in an exploratory search task are different. Using this knowledge, we were able to suggest roles for these users to follow for the rest of their search sessions. Using simulations on two different user study datasets for what would happen if such suggestions were followed, we were able to show promising results highlighting the benefit of role mining.

This work has a few limitations. For instance, it assumes that a user's role or in general search behavior can be extracted using his actions. The proposed algorithm also ignores any prior knowledge or preferences about the users. This limitation can be addressed by creating a hybrid model that takes into consideration a user's past behavior and preferences, in addition to current actions during a search session. For the future work, we also plan to derive role patterns through the users' search session relying on meta-roles, not necessarily defined within a role taxonomy.

7. ACKNOWLEDGEMENTS

Authors thank the UPS University and the IRIT laboratory for their grants. Authors also thank participants of both user studies, supported by the US Institute of Museum and Library Services (IMLS) Early Career Development grant # RE-04-12-0105-12.

8. REFERENCES

- [1] W. K. Edwards. Policies and roles in collaborative applications. In *ACM CSCW*, pages 11–20, 1996.
- [2] H. Feild and J. Allan. Task-aware query recommendation. In *ACM SIGIR*, pages 83–92, 2013.
- [3] C. Foley and A. F. Smeaton. Synchronous collaborative information retrieval: Techniques and evaluation. In *ECIR*, pages 42–53, 2009.
- [4] J. Freyne, B. Smyth, M. Coyle, E. Balfe, and P. Briggs. Further Experiments on Collaborative Ranking in Community-Based Web Search. *Artificial Intelligence Review*, 21:229–252, 2004.
- [5] S. A. Golder and J. Donath. Social roles in electronic communities. In *Proc. of Association of Internet Researchers Conference*, 2004.
- [6] G. Golovchinsky, O. Qvarfordt, and J. Pickens. Collaborative information seeking. *IEEE Computer*, 42(3):47–51, 2009.
- [7] K. Henderson, B. Gallagher, T. Eliassi-Rad, H. Tong, S. Basu, L. Akoglu, D. Koutra, C. Faloutsos, and L. Li. Rolx: Structural role extraction & mining in large graphs. In *ACM SIGKDD*, pages 1231–1239, 2012.
- [8] J. Hyldegard. Collaborative information behaviour - exploring Kuhlthau's Information Search Process model in a group-based educational setting. *Inf. Process. Manage.*, 42:276–298, 2006.
- [9] J. Hyldegard. Beyond the search process - exploring group members' information behavior in context. *Inf. Process. Manage.*, 45(1):142–158, 2009.
- [10] H. Joho, D. Hannah, and J. M. Jose. Revisiting ir techniques for collaborative search strategies. In *ECIR*, pages 66–77, 2009.
- [11] C. C. Kuhlthau. Keynote address: Towards collaboration between information seeking research and information retrieval. *Inf. Res.*, 10(2), 2005.
- [12] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *ACM WWW*, pages 591–600, 2010.
- [13] Y. Lv and C. Zhai. Positional relevance model for pseudo-relevance feedback. In *ACM SIGIR*, pages 579–586, 2010.
- [14] A. McCallum, X. Wang, and A. Corrada-Emmanuel. Topic and role discovery in social networks with experiments on enron and academic email. *J. Artif. Int. Res.*, 30(1):249–272, 2007.

- [15] M. R. Morris and E. Horvitz. Searchtogether: an interface for collaborative web search. In *UIST*, pages 3–12, 2007.
- [16] M. R. Morris, A. Paepcke, and T. Winograd. Teamsearch: Comparing techniques for co-present collaborative search of digital media. In *Tabletop 2006*, 2006.
- [17] M. R. Morris, J. Teevan, and S. Bush. Enhancing collaborative web search with personalization: groupization, smart splitting, and group hit-highlighting. In *CSCW*, pages 481–484, 2008.
- [18] K. Nowicki and T. A. Snijders. Estimation and prediction for stochastic blockstructure. *Journal of the American Statistical Association*, 96(455):1077–1087, 2001.
- [19] A. Pal and S. Counts. Identifying topical authorities in microblogs. In *ACM WSDM*, pages 45–54, 2011.
- [20] J. Pickens, G. Golovchinsky, C. Shah, P. Qvarfordt, and M. Back. Algorithmic mediation for collaborative exploratory search. In *ACM SIGIR*, pages 315–322, 2008.
- [21] T. Rodden. A Survey of CSCW Systems. *Interacting with Computers*, 3(3):319–353, 1991.
- [22] C. Shah and R. González-Ibáñez. Evaluating the synergic effect of collaboration in information seeking. In *ACM SIGIR*, pages 913–922, 2011.
- [23] C. Shah and R. Gonzalez-Ibanez. Coagmento - a system for supporting collaborative information seeking. In *ASIST*, pages 9–12, 2011.
- [24] C. Shah, J. Pickens, and G. Golovchinsky. Role-based results redistribution for collaborative information retrieval. *Inf. Process. Manage.*, 46(6):773–781, 2010.
- [25] B. Smyth, E. Balfe, P. Briggs, M. Coyle, and J. Freyne. Collaborative Web Search. In *IJCAI*, pages 1417–1419, Acapulco, Mexico, 2003.
- [26] L. Soulier, L. Tamine, and W. Bahsoun. A collaborative document ranking model for a multi-faceted search. In *Springer AIRS*, pages 109–120, 2013.
- [27] L. Soulier, L. Tamine, and W. Bahsoun. On domain expertise-based roles in collaborative information retrieval. *Information Processing & Management*, 2014, to appear.
- [28] S. Talja. Information sharing in academic communities: Types and levels of collaboration in information seeking and use. *New Review of Information Behaviour Research*, 3:143–160, 2002.
- [29] J. Teevan, S. T. Dumais, and E. Horvitz. Personalizing search via automated analysis of interests and activities. In *ACM SIGIR*, pages 449–456, 2005.
- [30] J. Xu and W. B. Croft. Query expansion using local and global document analysis. In *ACM SIGIR*, pages 4–11, 1996.