

Pseudo-Collaboration as a Method to Perform Selective Algorithmic Mediation in Collaborative IR Systems

Roberto González-Ibáñez^{1&2}, Chirag Shah¹, Ryan W. White³

¹ School of Communication & Information, Rutgers University, New Brunswick, NJ 08901 USA

² Departamento de Ingeniería Informática, Universidad de Santiago de Chile, Santiago, Chile

³ Microsoft Research, Redmond, WA 98052 USA

{rgonzal, chirags}@rutgers.edu; ryanw@microsoft.com

ABSTRACT

Traditional recommendation systems suggest results based on data collected from users' actions. Many of the newer information retrieval (IR) systems incorporate social search or collective search signals as an extension to standard term-based retrieval algorithms. Systems based on social or collaborative search methods, however, do not consider when, how, and to what extent such support could help or hurt their users' search performance. In this poster we propose a novel approach of selective algorithmic mediation capable of identifying when a user should be aided by a collaborator and to what extent such help could enhance search success. We demonstrate the applicability and benefits of our approach through simulations using a pseudo-collaboration method on the log data of individual users and pairs of users gathered during a laboratory study with 131 participants. The results show that our approach can improve the search performance of both individual searchers and others collaborating intentionally by identifying and recommending regions in search processes with best chance of improvements, thus increasing the likelihood that users find more useful information with less effort.

Keywords

Collaborative information retrieval; Search effectiveness.

INTRODUCTION

Collaboration is often considered to be an expected or necessary solution for complex problems (Denning, 2007), including those relating to search (Twidale et al., 1997). However, it is often not known if and when someone could benefit from collaborating in retrieval scenarios. The information retrieval (IR) community has developed a number of methods to help improve a user's search processes, including query suggestions, results recommendation, information filtering, and personalization. Collaboration is another such possibility that could enhance a user's information seeking by obtaining richer and more diverse information, as well as facilitating social connections and learning (Shah & González-Ibáñez, 2011; Twidale et al.,

1997). We are interested in predicting and evaluating the feasibility of collaboration in a non-collaborative situation, along with projected benefit of such collaboration.

Research on collaborative IR (CIR) is often characterized by the nature of mediation in collaboration (Pickens et al., 2008). During *system* or *algorithmic mediation*, the system acts as an active agent and provides mediation among the collaborators to enhance their productivity and experience. System-mediated collaboration can improve productivity in search tasks, but it often assumes pre-defined roles of the users with little or no flexibility offered to the collaborators during the process. Conversely, in *user* or *interface mediation* the control lies with the collaborators, with the system being a passive actor. In these situations, the users drive the collaboration, and the system primarily provides a range of interface and algorithmic functionality. User-mediated collaborations, while providing greater flexibility and control to the users, are typically limited by what participants know, do, and agree on without getting active system assistance. That said, even when partial support from the system is offered, it often goes unused (Pickens et al., 2008).

To address these shortcomings, we propose a method that leverages *pseudo-collaboration* to perform *selective* algorithmic mediation in CIR systems. Pseudo-collaboration is a type of CIR. Unlike similar methods such as collaborative filtering (Herlocker et al., 2004), pseudo-collaboration is intended to determine if an individual searching for information could benefit from collaborating with someone else. If so, the method informs when and how such collaboration should take place. Pseudo-collaboration uses simulations capable of projecting the search process of users as they were working with others. Such simulations allow the CIR to evaluate the benefit and cost of aiding a user at different stages of their search process using search sessions of other users performing the same or similar tasks.

Pseudo-collaboration lets users retain control over their search processes, with the system making recommendations when they are beneficial. If users ignore or dismiss the suggestions, the search remains user mediated, and if the user accepts them, it could become system-mediated.

In this poster we describe pseudo-collaboration in more detail and present results from an evaluation of the method using log data from a laboratory study of CIR.

BACKGROUND

Before proceeding, it is important to situate our research with respect to existing work on collaboration and discuss mediation in CIR in more detail.

Contextual Definition

As described above, pseudo-collaboration is a type of CIR. Unlike related techniques such as collaborative filtering, pseudo-collaboration focuses not only on the system (algorithmically-generating recommendations from other users) but also on the user perspective (fostering collaboration between users). Rather than simply providing suggestions based on aggregated results, pseudo-collaboration provides a mechanism to project the search process of a user and identify when and how collaboration (either implicit/unintentional or explicit/intentional) would result in a benefit to the user. Figure 1 depicts pseudo-collaboration with respect to CIR and collaborative filtering.

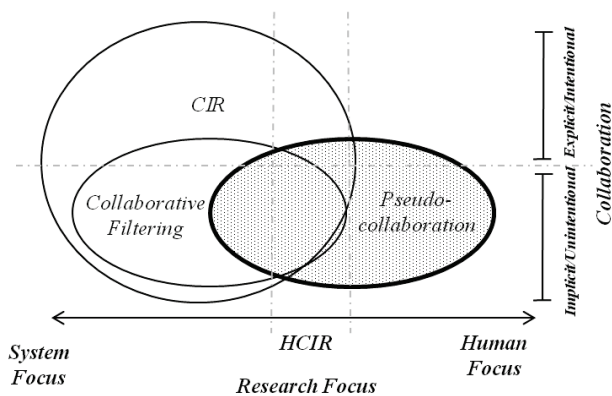


Figure 1: A conceptual depiction of pseudo-collaboration with respect to CIR and collaborative filtering.

Mediation in Collaborative IR

Research on collaborative IR is often characterized by the nature of mediation in collaboration (Pickens et al., 2008):

- *System or algorithmically mediated.* Here, the system acts as an active agent and provides mediation among the collaborators to enhance their productivity and experience. A recent example under this category is Querium (Golovchinsky et al., 2012).
- *User or interface mediated.* Here, the control lies with the collaborators, with the system being a passive component. The users drive the collaboration, and system primarily provides various functions on the interface level. Examples include SearchTogether (Morris & Horvitz, 2007), and Coagmento (Shah, 2010).

Different researchers have shown how system-mediated collaboration could improve productivity in search tasks (e.g., (Pickens et al., 2008), (Shah, 2010)), but they often assume pre-defined roles of the users with little or no flexibility offered to the collaborators during the process. For instance, Shah et al. (2010) showed how a system could help two collaborators playing different roles achieve re-

trieval results that are more relevant and novel than what either of them could have working individually. However, this setting could only work if the users already know their roles, responsibilities, and abilities, and do not change them throughout the collaborative process. On the other hand, user-mediated collaborations, while providing greater flexibility and control to the users are typically limited by what the individuals involved in collaboration know, do, and agree on without getting active assistance from the system. Even when partial support from the system is incorporated, it often goes unused. For instance, Morris and Horvitz (2007) found that “split search” feature of their SearchTogether system was underutilized even though it could have helped users perform more effective division of labor.

As mentioned in the previous section, to overcome shortcomings of both the approaches for mediating collaboration, we propose pseudo-collaboration, a unique method that performs selective mediation for CIR.

PSEUDO-COLLABORATION

In general terms, pseudo-collaboration is intended to aid single users’ search processes by simulating, assessing, and selecting collaboration with other users. The aim of pseudo-collaboration is to increase the probability of finding useful information at a low cost. We refer to this approach as *pseudo-collaboration* because of the lack of explicitness and intentionality in which convenient and impersonal combinations of users’ search sessions take place. In order to optimize pseudo-collaboration, we should consider some of the implications of teaming a user up with someone else. For example, assume that users A and B do not know each other, but they have common information needs. On day 1, user A completes a search session. Then on day 2, user B starts his session. Pseudo-collaboration works through four steps: (1) Identify that A and B have similar information needs. (2) Predict what would happen if B would be aided with what A did in his search session. (3) Determine when the benefit would be significant. (4) Select elements from A’s search behavior such as queries and information encountered that could help improve B’s search performance.

While each of the steps indicated above are necessary and challenging, this poster focuses mainly on the second and third steps. The first step, on the other hand, is a challenging research topic in itself and as we explain later, we skipped the evaluation of topic similarity in our experiments since the data we used belongs to users performing the same task. Regarding the fourth step, we considered only Web pages and queries for evaluation purposes.

Search Process Projection

As a first step in building our pseudo-collaboration method we projected in time the search session of a user in order to predict what would happen if they were teamed up with someone else (so called *search process projection*). Search projection is carried out through simulations consisting of

search session alignment based on topic similarity. If two or more sessions are found to be similar based on a given criteria (e.g. query similarity over the first minutes of a search session), then all possible combinations are generated by projecting the search process as long as the individual search sessions last. Then such combinations are compared at different times selecting only those that produce significant improvement compared to individual search.

We simplified the actions of users in online search into three major stages, namely: query formulation, search engine result page (SERP) exploration, and content evaluation. Users' actions may comprise *search trails* involving query reformulation, browsing, and finding useful and/or relevant material. Each search trail may require from few to several actions that can be mapped to any of these states. Pseudo-collaboration in this regard generates optimal combinations of users' search processes by merging users' actions in a way that the search performance is maximized along the search process, with the aim of providing pertinent correct and timely assistance to searchers.

EVALUATION

We evaluated pseudo-collaboration using the search sessions of 11 individual participants and 120 collaborative participants who performed an exploratory search task as part of a large study by Shah and González-Ibáñez (2011). All participants performed the following task:

“A leading newspaper has hired your team to create a comprehensive report on the causes, effects, and consequences of the recent gulf oil spill. As a part of your contract, you are required to collect all the relevant information from any available online sources that you can find.

To prepare this report, search and visit any website that you want and look for specific aspects as given in the guideline below. As you find useful information, highlight and save relevant snippets. Make sure you also rate a snippet to help you in ranking them based on their quality and usefulness. Later, you can use these snippets to compile your report, no longer than 200 lines, as instructed.

Your report on this topic should address the following issues: description of how the oil spill took place, reactions by BP as well as various government and other agencies, impact on economy and life (people and animals) in the gulf, attempts to fix the leaking well and to clean the waters, long-term implications and lessons learned.”

Since they all performed the same task, we could be sure that there was consistency in their information needs (Step 1 from earlier). The search logs for each user contained data from 20 minutes of active search, collection, and evaluation (relevance judgment) of information collected to accomplish the task. To address Steps 2 and 3 outlined earlier, we performed an exhaustive search for optimal pairs of participants by generating all possible combinations of participants' search sessions (130 per user, which resulted in 8,515 pairs for all 131 participants in the study).

To compare users and the user pairs produced by our method, we defined two measures – *effectiveness* and *efficiency* – that were computed on-the-fly during the search process. Effectiveness was defined as the precision of a given user or pair (u) in finding useful pages at time t , in terms of the ratio between their useful coverage (*UsefulCoverage*) and overall coverage (a count of all distinct pages visited) at time t (Eq. 1). *UsefulCoverage* is the number of distinct pages visited by u for at least 30 seconds, suggesting satisfaction with their content (Fox et al., 2005; Shah & González-Ibáñez, 2011; White & Huang, 2010). We used this dwell time threshold to estimate the utility of pages automatically, without requiring explicit judgments. Note that this implicit measure of usefulness maps to positive relevance labels in our dataset; 70% of the pages found useful according to the dwell-time threshold were also relevant according to participants' explicit judgments.

$$Effectiveness(u, t) = \frac{UsefulCoverage(u, t)}{OverallCoverage(u, t)} \quad (1)$$

We also measured search performance in terms of the efficiency of u as the ratio between effectiveness and the number of queries that need to be formulated (cost) to find useful pages (Eq. 2).

$$Efficiency(u, t) = \frac{Effectiveness(u, t)}{NumQueries(u, t)} \quad (2)$$

For each user and each generated pair, we computed the above measures ($f(t)$) at each minute in the session and cumulatively for time slices from the start of the session (t_0) to the current time (t_c). Given the sequences of discrete time points, we computed the areas under the curve (AUC) for each time slice using the trapezoidal rule for non-uniform grids as a numerical method of integration (Eq. 3).

$$\int_{t_0}^{t_c} f(t) dt \approx \frac{1}{2} \sum_{k=t_0}^{t_c} (t_{k+1} - t_k) (f(t_{k+1}) + f(t_k)) \quad (3)$$

We then compared the AUCs for each measure with those of the generated pairs (130 per user) at different intervals during the session selecting only those that produced significant improvements at $p < .05$ (measured using the z-score since Shapiro-Wilks tests showed that our data were normally distributed), thus increasing the likelihood of useful information encounters at a lower cost.

RESULTS AND DISCUSSION

Figure 2 depicts three views of the performance of pseudo-collaboration. Figure 2a shows the average scores and Figure 2b the cumulative AUC for each measure across the session for users as individual units versus pseudo-collaborative pairs (Figure 2a). It is clear that our method can increase the likelihood that a user finds more useful information with less effort (Figure 2b). The gains over individual users also increase as the session proceeds. We also found that although pseudo-collaboration could improve search performance throughout the search process, the fraction of collaborations that yield benefit drops to

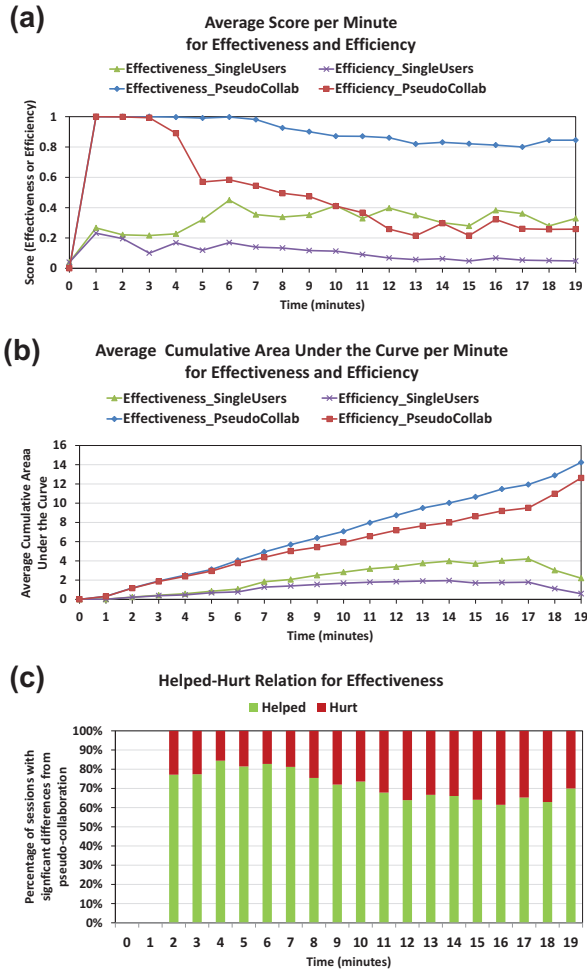


Figure 2. Three views of the performance of pseudo-collaboration. (a) Average scores in each minute, (b) Average AUC in each minute, and (c) Helped-hurt relation in each minute. Error bars are not displayed in (a) and (b) to avoid crowding the figure.

around 60% in advanced stages of the task (Figure 2c), perhaps because that is when searchers' information needs become more specific (a hypothesis to be tested in future work). No values are shown in Figure 2c for the first two minutes since there we did not observe any instances of significant benefit or harm at those points. Finally, in a comparison between real collaborative pairs and pseudo-collaboration, we found that all study participants being aided by our method could outperform the search performance obtained by working with their actual collaborator. Also, an analysis of one-minute segments of the search sessions in this comparison revealed that 21% of them were more effective and efficient with pseudo-collaboration, although only around 2% were significantly better at $p < .05$ (again using the z-score).

CONCLUSIONS

Using pseudo-collaboration we showed how to identify points or regions of an individual's search process where

they could benefit from collaborating with another searcher with the same/similar information needs. While it is challenging to ensure the success of a collaborative IR project, we demonstrated how we could use measures like *effectiveness* and *efficiency* to estimate the benefit to a searcher if they accept a recommendation for collaboration. Pseudo-collaboration, therefore, offers a solution for selective collaboration that allows one to dial between user-mediated and system-mediated CIR. For future study, we would like to evaluate the effectiveness of our approach as part of a comparison of collaborating searchers. We believe that in scenarios where multiple users are performing similar or identical tasks at the same time (something that may be common both on the Web or in large enterprise settings), pseudo-collaboration could lead users to actual collaboration based on the results of the search projection process using past data from other users who performed similar tasks (e.g., high projection overlap between concurrent searchers may indicate a collaboration opportunity).

Our next step is to carry out an experiment to evaluate pseudo-collaboration with data collected from a large-scale search engine. Unlike the study presented here, we will need to address the first step of pseudo-collaboration (i.e., identify topic similarity). We plan to do this with heuristics or methods already established in this research domain.

REFERENCES

- Denning, P.J. (2007). Mastering the mess. *CACM*, 50(4): 21–25.
- Fox, S., Karnawat, K., Mydland, M., Dumais, S. & White, T. (2005). Evaluating implicit measures to improve web search. *ACM TOIS*, 23(2): 147–168.
- Golovchinsky, G., Dunnigan, A. & Diriyeh, A. (2012). Designing a tool for exploratory information seeking. *Proc. CHI*, 1799–1804.
- Herlocker, J.L., Konstan, J.A., Terveen, L.G. & Riedl, J.T. (2004). Evaluating collaborative filtering recommender systems. *ACM TOIS*, 22(1): 5–53.
- Morris, M. R., & Horvitz, E. (2007). SearchTogether: An interface for collaborative web search. *Proc. UIST*, 3-12.
- Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P. & Back, M. (2008). Algorithmic mediation for collaborative exploratory search. *Proc. SIGIR*, 315–322.
- Shah, C. (2010). Coagmento: a collaborative information seeking, synthesis and sense-making framework. *CSCW*.
- Shah, C., Pickens, J. & Golovchinsky, G. (2010). Role-based results redistribution for collaborative information retrieval. *IP&M*, 46(6): 773–781.
- Shah, C. & González-Ibáñez, R. (2011). Evaluating the synergic effect of collaboration in information seeking. *Proc. SIGIR*, 913–922.
- Twidale, M., Nichols, D. & Paice, C. (1997). Browsing is a Collaborative Process. *IP&M*, 33(6): 761–783.
- White, R.W. & Huang, J. (2010). Assessing the scenic route: Measuring the value of search trails in web logs. *Proc. SIGIR*, 587–594.