Introduction to the Special Issue

Gene Golovchinsky
FX Palo Alto Laboratory, Inc.
Palo Alto, CA, USA

gene@fxpal.com

Meredith Ringel Morris
Microsoft Research
Redmond, WA, USA

merrie@microsoft.com

Jeremy Pickens
FX Palo Alto Laboratory, Inc.
Palo Alto, CA, USA

jeremy@fxpal.com

This special issue brings together papers that describe some of the many ways that collaborative information seeking manifests itself. Some papers report on collaborative practices in a range of domains, including medical (Hertzum), legal (Attfield *et al.*), and online Q&A (Gazan). Others propose and evaluate models of collaborative activity (Evans and Chi; Evans *et al.*; Wilson and schraefel; Foley and Smeaton), and others describe systems and algorithms that support collaboration in various ways (Boydell and Smyth; Fernandez-Luna *et al.*, Halvey *et al.*, Morris *et al.*; Shah *et al*.).

Our perspective is that the high-level goal of this research is to improve the way that people manage their collaborative information seeking online. Thus we study how people work to help us understand the range of problems and opportunities, we build models and theories to make sense of the observed phenomena, and we build and evaluate systems to test our understanding.

In this context, it is worthwhile to consider the notion of computer-mediated collaborative information seeking broadly, before diving into the details. The term "collaborative search" has been used to describe a wide range of people's behaviors and the computer systems used to support them. Here we will focus on only those situations when the collaboration was mediated in some important way by computer-based tools. Thus we consider the use of generic communication tools, recommendation systems, Q&A systems, social search that leverages social networks for information seeking, co-browsing and link sharing, and more complex CSCW systems. While these situations all differ in important ways -- important for those who use them and for those who design for them -- they also have significant commonalities (Pickens and Golovchinsky, 2007; Golovchinsky *et al*., 2009).

Foremost among the commonalities is the presence of multiple people engaged in a variety of information seeking activities. The activities consist of users' behavior and of systems with common characteristics. With respect to behavior, we can classify people's activity based on the notion of *collaborative intent*; with respect to system characteristics, we can classify information seeking systems in terms of their ability to represent individual people's actions and data.

When people work together, do they have explicitly or implicitly shared information needs? A recommender system is a good example of implicit sharing: a searcher who receives a recommendation based on similar behaviors or opinions (traits) (Teevan *et al.*, 2009) of other

people may not in fact have the same information need in mind. The success of such a system relies then on its ability to *infer* similar information needs from similar behaviors or opinions. Compare this with SearchTogether (Morris and Horvitz, 2007), a system that allows groups of people to share queries and search results in pursuit of a common information need (task). The *explicit* decision to work together on a search task frees the system to concentrate on managing the sharing of information.

Collaborative information seeking support systems can also be classified by the degree to which the system represents each person's information independently with respect to search activity. In a social search environment such as Aardvark.com, for example, mediation is limited to communication among people to articulate an information need; the rest of the process takes place without any mediation, and often outside the system that manages the communication. SearchTogether, on the other hand, allows people to exchange queries and search results through a shared user interface, but all queries are passed through the same system without regard to who issued each query. Thus we can characterize this kind of collaborative search as having UI-level mediation. A recommender system, on the other hand, keeps track of who did what, and aggregates that data to suggest information to the next searcher. This requires mediation on the retrieval algorithm level. We call this continuum -- communication, UI-level mediation, algorithmic mediation -- the *depth of mediation* dimension.

We can also look at collaboration in terms of information flow among collaborators. We distinguish three possible situations:

1. The degenerate case: No information flows among collaborators other than manually-exchanged results.
2. The asymmetric case: Information can flow only in certain ways, from some collaborators to others, but not vice versa. For example, the behavior of the person receiving recommendations aggregated from the behaviors of earlier searchers does not affect those earlier searchers in any material way with respect to the information need that had caused them to seek information.
3. The symmetric case: Finally, systems such as SearchTogether or Cerchiamo (Pickens *et al.*, 2008) allow influence to flow among all participants. While participants' roles may affect what information they are shown by the search system at any particular time, in principle all information is available to all collaborating searchers.

We do not mean to imply that people have to work synchronously in the tightly-coupled manner characteristic of some groupware systems. Rather, what is important is that the data that is presented to collaborating searchers be synchronized across collaborating searchers' client software. Searchers can use the systems at their pace, use different interfaces, or work at different times, but still share the same consistent underlying data with respect to a particular search task. We therefore call this dimension *data synchronization*.

Returning to the papers in this issue, we see that Attfield *et al.* and Hertzum describe instances of explicit collaboration with communication as the only form of mediation. Gazan looks at a

combination of collaborative filtering (implicit collaboration) with episodes of explicit problem-solving collaboration for which only communication is mediated. Evans *et al*. assess the effectiveness of communication with others for information seeking vs. individual information seeking, and Evans and Chi examine a range of implicit and explicit collaboration. Morris *et al.*, Wilson and schraefel, and Halvey *et al.* all describe UI-level mediation; while the first two papers focus on explicit collaboration, Halvey *et al.*'s system suggests a transition between the two. Finally, Luna *et al.*, Shah *et al.*, and Foley and Smeaton all describe systems and algorithms for supporting explicit, algorithmically-mediated collaboration, while Boydell and Smyth focus more on implicit collaboration. Figure 1 shows how the papers cluster in intent vs. mediation.

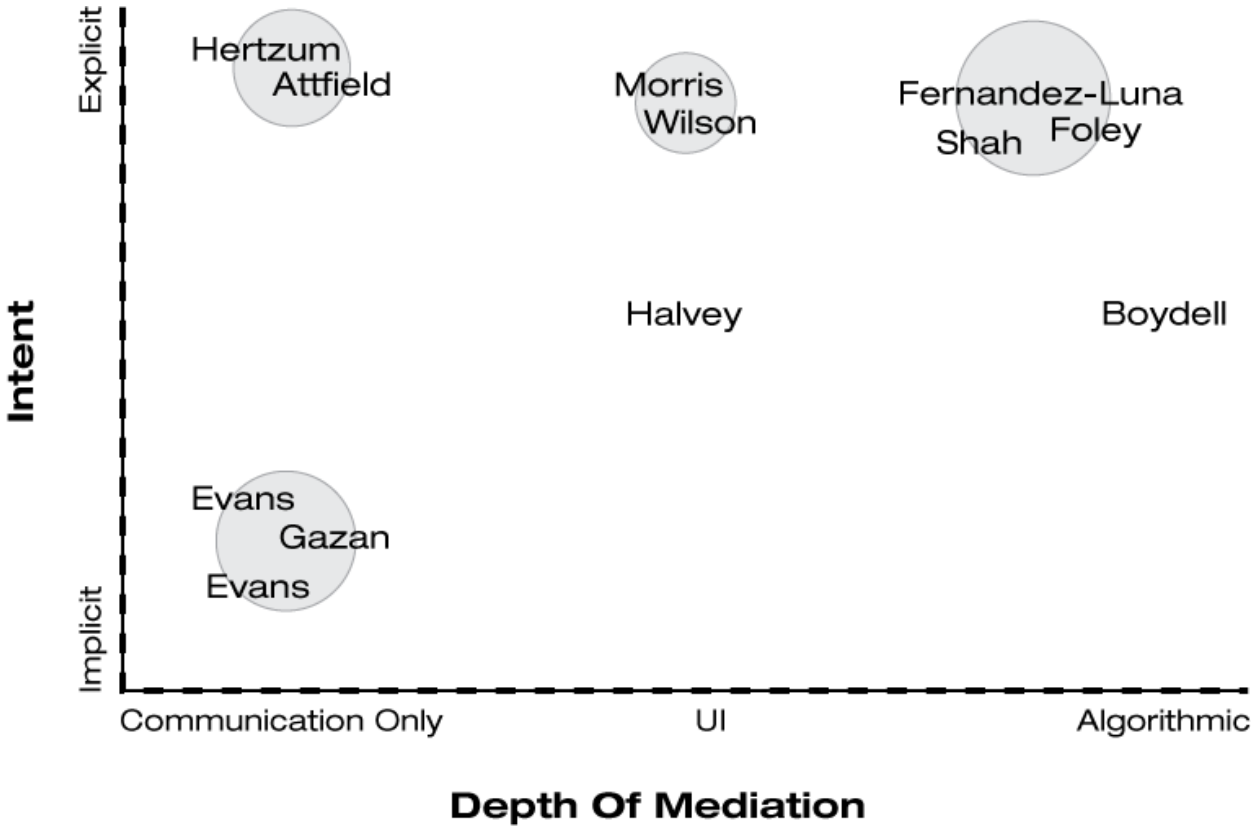

Figure 1. Intent vs. Depth of mediation for a given task. Grey circles represent papers that are in the same point in our design space. Papers are identified by the last name of the first author.

Similarly, we can classify papers based on intent and data synchronization, as shown in Figure 2. We see three clusters – Hertzum and Attfield *et al.* in the explicit-nonsynchronized quadrant; Fernandez-Luna *et al.*, Foley and Smeaton, Morris et al., Shah *et al.*, and Wilson and schraefel in the explicit-synchronized quadrant; and Evans *et al.*, Evans and Chi, and Gazan in the implicit, non-synchronized quadrant. Boydell and Smyth occupy the middle ground between implicit and explicit collaboration with task-related data synchronization, while Halvey *et al.*, describe a similar approach with synchronized data.

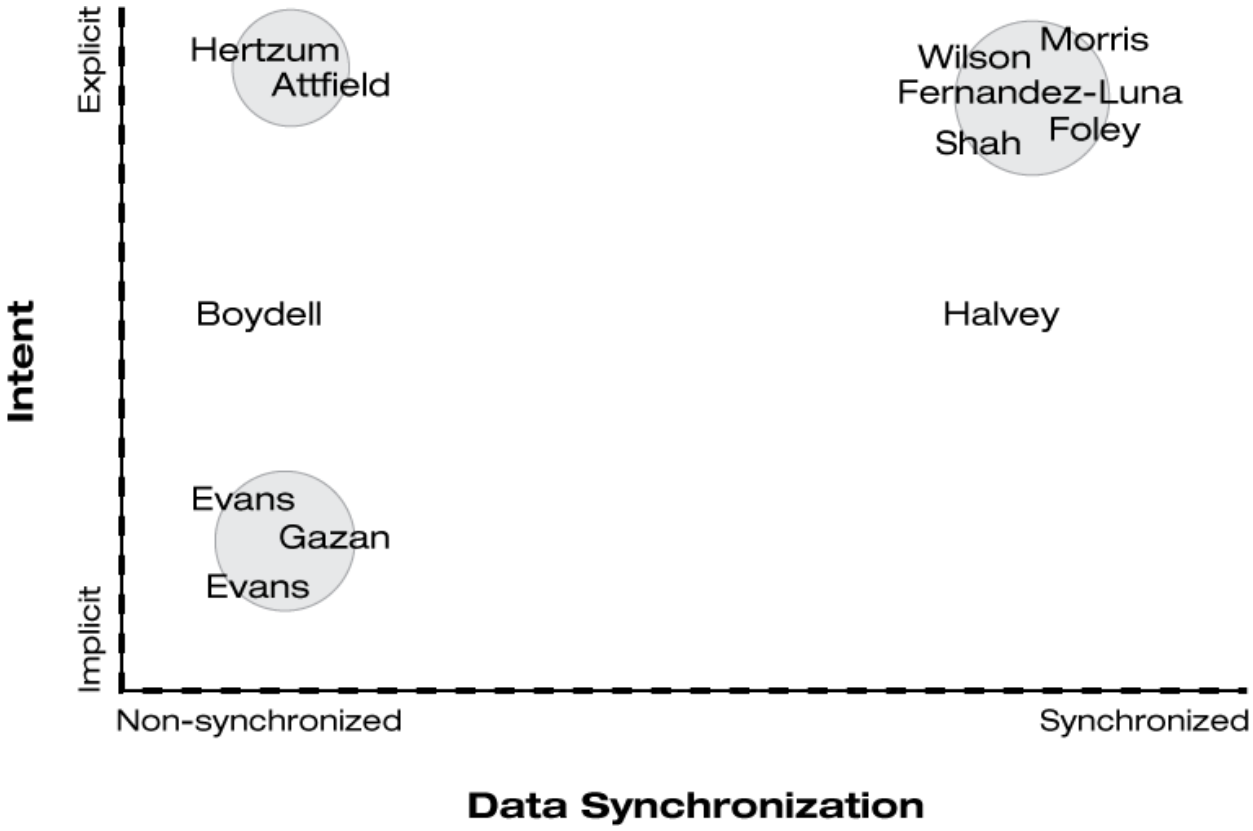

Figure 2. Intent vs. data synchronization for a given task. Grey circles represent papers that are in the same point in our design space. Papers are identified by the last name of the first author.

.

The papers in this issue represent a broad range of interpretations and approaches to collaborative search, but they are just the beginning. We expect that a better understanding of collaboration will lead to more system building, which in turn will expose both users and designers to more opportunities to establish novel patterns of work, producing more evaluations, and improving the richness of our models. As this body of work grows, existing challenges such as evaluation metrics and test collections will also need to be addressed.

References

Attfield, S. J., Blanford, A.E., and Makri, S. (this issue) Social and Interactional Practices for Disseminating Current Awareness Information in an Organisational Setting.

Boydell, O., and Smyth, B. (this issue) Social Summarization in Collaborative Web Search

Evans, B.M. and Chi. E.H. (this issue) An Elaborated Model of Social Search.

Evans, B.M., Kairam, S. and Pirolli, P. (this issue) Do Your Friends Make You Smarter?: An Analysis of Social Strategies in Online Information Seeking.

Fernandez-Luna, J.M., Huete, J.F., Perez-Vazquez, R., and Rodriguez-Cano, J.C. (this issue) CIRLab: A Groupware Framework for Collaborative Information Retrieval.

Foley, C. and Smeaton, A.F. (this issue) Division of Labour and Sharing of Knowledge for Synchronous Collaborative Information Retrieval.

Gazan, R. (this issue) Microcollaborations in a Social Q&A Community.

Golovchinsky, G., Qvarfordt, P., and Pickens, J. (2009) Collaborative Information Seeking. IEEE Computer, 42(3). March 3, 2009

Halvey, M., Vallet, D., Hannah, D., Feng, Y., and Jose, J. (this issue) An Asynchronous Collaborative Search System for Online Video Search.

Hertzum, M. (this issue) Breakdowns in Collaborative Information Seeking: A Study of the Medication Process.

Morris, M.R., Fisher, D. and Wigdor, D. (this issue) Search on Surfaces: Exploring the Potential of Interactive Tabletops for Collaborative Search Tasks.

Morris, M.R. and Horvitz, E. (2007) SearchTogether: An Interface for Collaborative Web Search. *Proceedings of ACM UIST 2007*, 3 - 12

Pickens, J. and Golovchinsky, G. (2007) Collaborative Exploratory Search. *Proceedings of the Human Computer Information Retrieval (HCIR) Workshop*, Nov 2, 2007.

Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P., and Back, M. (2008) Algorithmic Mediation for Collaborative Exploratory Search. *Proceedings of SIGIR 2008*, 315-322.

Shah, C., Pickens, J., and Golovchinsky, G. (this issue) Role-Based Results Redistribution for Collaborative Information Retrieval.

Teevan, J., Morris, M.R., and Bush, S. (2009) Discovering and Using Groups to Improve Personalized Search. *Proceedings of WSDM 2009*, 15-24.

Wilson, M.L. and schraefel, .mc. (this issue) Evaluating Collaborative Information Seeking Interfaces with a Search-Oriented Inspection Method and Re-framed Information Seeking Theory.