

Really Supporting Information Seeking: A Position Paper

Nicholas J. Belkin
SCILS, Rutgers University
4 Huntington Street
New Brunswick, NJ 08901, USA
+1 732 932 7500 x8271
nick@belkin.rutgers.edu

1. Introduction

It has been quite clear for quite some time that people engage in a wide variety of interactions with information when engaging in information seeking, whether in multiple information seeking episodes over time, or during the course of a single information seeking episode. Yet systems, both experimental and operational, intended to support information seeking, have been limited, with few exceptions, to the support of only one kind of information seeking strategy (ISS), and only one kind of interaction with information, non-contextualized specified search for particular information objects. In this paper, I argue that truly effective systems for supporting information seeking must be able to support the variety of interactions with information that people engage in, with knowledge of the context and situation which led them to engage in information seeking behavior, and, in the best case, without having to leave their task environment at all. I do this by indicating what I feel are significant challenges to which research in this area should respond.

2. Information-related goals, tasks and intentions

There is substantial and convincing evidence that the goals that lead people to engage in information behavior, the tasks associated with those goals, and with their behaviors, and the intentions underlying the behaviors, substantially affect their judgments of usefulness of information objects, and the ways in which they interact (or would wish to interact) with information objects. The challenges here lie in three spheres.

First, in the ability to characterize and differentiate among information-related goals, tasks and intentions in some principled manners that go beyond straightforward listings, that will apply across a wide variety of contexts, and from which design principles for IR systems can be inferred.

Second, we need to develop methods for inferring information-related goals, tasks and intentions from implicit sources of evidence, such as previous or concurrent behaviors. Without such techniques, any characterization of these factors, or even specification of how to respond to them in system design, is fruitless.

Third, going from characterization and identification of goals, tasks and intentions, to IR techniques which actually respond effectively to them, is a challenge that has been barely noted, much less addressed, to date.

3. Understanding and supporting information behaviors other than specified search

People engage in a wide variety of types of interactions with information, of which specified searching, as represented by, e.g., normal Web search engines, and standard IR models, is only one among many, and perhaps not the most important. For instance, since we know, from a good number of different theoretical and empirical studies, that people have substantial difficulties in specifying what it is (that is, what information objects) that would help them to realize their goals, only considering specified search as the basis for IR models and techniques, is clearly inadequate, and inappropriate. Furthermore, a number of studies have demonstrated that people do engage in a variety of different information behaviors within even a single information seeking episode. However, there is still little known about the nature, variety and relationships of different information behaviors, and the situations that lead people to engage in any one behavior, or sequence of behaviors. Without this basic knowledge, there is little hope for the development of systems which can proactively support multiple behaviors.

4. Characterizing context

The challenges here are obvious, and multiple. We need to have theories and data which lead not only to identification of at least classes of contextual features, but which also tell us in what ways these features do in fact affect information behaviors. We need to have ways to identify appropriate contextual features without explicitly eliciting them. We need to develop IR system techniques which can actually take account of knowledge of these features, individually and in combination, and we need to develop experimental methods and evaluation measures which will allow comparison of the effect of contextually-based interventions.

5. Personalization

The challenge with respect to personalization is first to consider the dimensions along which personalization could and should take place; then to identify the factors or values in each dimension that would inform personalization on that dimension; then to investigate how the different factors or types of evidence interact with one another; and finally, to devise techniques which take account of these results in order to lead to a really personalized experience.

6. Integration into the task environment

Engaging in information interaction in an IR system is almost always a consequence of a person's wanting to achieve some other goal, or accomplish some other task. In this sense, such engagement is inevitably a distraction from that other task. So, we might say that an ultimate goal of, and challenge for IR research is to arrange things such that a person never has to engage in interaction with a separate IR system at all (although I'm quite willing to agree that there are certainly circumstances in which such engagement might indeed be desirable).

7. Other important challenges (in brief)

Taking account of affect;

New evaluation paradigms for interactive information retrieval;

(In)Formal models of interactive information retrieval.