

Intelligent Information Retrieval: Whose Intelligence?

Nicholas J. Belkin*

School of Communication, Information and Library Studies
Rutgers University
4 Huntington Street
New Brunswick, NJ 08901-1071 USA
+1 908 932-8585
belkin@scils.rutgers.edu

Contents

1. What could we mean by intelligent information retrieval?
2. What constitutes an information retrieval system?
3. Where could intelligence be manifest in an information retrieval system?
4. Information retrieval systems as support for interaction with information.

Abstract

The concept of 'intelligent' information retrieval was first mooted in the late 1970s, but had lost currency within the information retrieval community by at least the early 1990s. With the popularity of the concept of 'intelligent agents', it appears that the idea of intelligent information retrieval is again in general vogue. In this paper, I attempt to show that the naive concept of intelligent information retrieval, based on the idea of agency, misses the essence of intelligence in the information retrieval system, and will inevitably lead to dysfunctional information retrieval. As a counter-proposal, I suggest that true intelligence in information retrieval resides in appropriate allocation of responsibility amongst all the actors in the information retrieval system, and that intelligent information retrieval will be achieved through effective support of people in their various interactions with information.

1. What could we mean by intelligent information retrieval?

Intelligent information retrieval (IR) has been variously defined by different people, but a consistent theme has been one of the machine (or program) doing something *for* the user, or the machine (or program) *taking over* some functions that previously had to be performed by humans (either user or intermediary). So, for instance, for Belkin et al. 87, an intelligent IR system was one in which the functions of the human intermediary were performed by a program, interacting with the human user. For Maes 94, on the other hand, intelligent IR is performed by a computer program (a so-called intelligent agent), which, acting on (perhaps minimal or even no explicit) instructions from a human user, retrieves and presents information to the user without any other interaction. Croft 87, introducing the *Information Processing and Management* Special Issue on Artificial Intelligence and Information Retrieval, suggests that intelligent IR is just *good* IR, meaning that it is inappropriate to ascribe intelligence to computer programs, and also meaning that good IR is that in which the programs (i.e. the representation, comparison and interaction methods implemented in the system) result in effective performance.

In all of these constructions, there is some idea that the intelligence (or goodness) of an intelligent IR system resides in the *built* system. Indeed, it is assumed by almost all people who have commented on intelligent IR, that the IR system is *only* the built system. In this paper, following Belkin et al. 83 and Belkin 93, I will argue that this assumption is, in itself, erroneous, that it has led to an inappropriate conception of what constitutes IR, and in particular, for

* Currently Fulbright Fellow, Department of Information Studies, University of Tampere, P.O. Box 607, FIN-3101 Tampere, tel: +358 31 215 7039, email lonibe@kielo.uta.fi

our purposes, has misconstrued the nature of intelligence in the IR system. Furthermore, following Bates 90, Belkin and Vickery 85, and Ingwersen 92 I will suggest both that interaction with information is the key to understanding intelligence in IR, and also that a fundamental problem for intelligent IR is how much, and what kind of support to offer users for such interaction.

2. What constitutes an information retrieval system?

IR systems are often construed as some collection of components and processes which take input, as a query to the system and as texts or information collected by the system, represent those inputs, compare them, and produce, as output, some set of texts or information objects predicted to be responsive to the query. I, and many others, have argued that it is a mistake to consider the IR system as only the built system. For extended argument on this issue, see Belkin 84. The crux of the issue is that, in order properly to consider all that happens in IR, and in order properly to evaluate the performance of IR systems, the boundary of the IR system must be outside the user, including the user within the IR system. From this point of view, the IR system consists of three major components:

the *user* in the system;

the *knowledge resource* to which the user has access and with which s/he interacts; and,

some person(s) and/or device(s) which supports and mediates the user's interactions with the knowledge resource (the *intermediary*).

The processes which are then considered to be significant within most conceptions of the IR system are:

representation (of user's information problem, of texts in the knowledge resource: e.g. indexing);

comparison (of representations of information problem and texts: e.g. retrieval techniques);

interaction (between user and intermediary: e.g. reference interview or human-computer interaction); and, sometimes,

judgment (of appropriateness of text to information problem, by the user: e.g. relevance judgments); and

modification (of representation of information problem: e.g. relevance feedback or query reformulation).

Belkin 84, Belkin et al. 83 and Ingwersen 92, for instance, have elaborated this general model of the IR system in several ways, primarily through being more explicit and detailed about the components, by extending the range of the different processes, by specifying how the processes are carried out, and especially by considering the nature of interaction in the IR system in much more detail. In particular, Ingwersen and Wormell 86 have suggested that it is appropriate to consider the direct interaction between user and text as a separate form of interaction, and this idea has been taken up in a more fundamental sense by Belkin 93, who suggests that IR should be considered explicitly as a form of interaction with information.

But in general, such elaborations have tended merely to extend somewhat the range of the IR system, without significantly changing the concept of the IR system itself. Thus, under this type of model, we understand the IR system to be composed of the components of user, intermediary and knowledge resource, related to one another by the processes of representation, comparison, interaction, judgment and modification. The significant point in this model, for

our purposes here, is that interaction, judgment and modification are inherent aspects of the IR system, and that the user is an inherent component of that system, not just some entity outside it, giving input and evaluating output.

3. Where could intelligence be manifest in an information retrieval system?

The inclusion of the user in the IR system, and the incorporation of interaction as a major process in IR, have some significant implications for how we might consider what would constitute intelligence in an IR system. For instance, under this view, the idea of the 'intelligent agent' seems untenable, at least in its most straightforward sense. That is, a program which takes a query as input, and returns documents as output, without affording the opportunity for judgment, modification and especially interaction with text, or with the program, is one which would not qualify as an IR system at all. In particular, such a program would fail to know about the user's information problem (relying only upon the query, some poor representation of that problem), and would fail to incorporate that one process which is known to improve retrieval performance significantly, interaction (especially, but not exclusively, through relevance feedback). So, although we might say that the representation and comparison processes might be performed well, and even 'intelligently', the system as a whole would not perform intelligently (if by that, we mean well, or effectively).

Another point which this view of the IR system raises is that there are some processes in the IR system which cannot be performed by any other component than the user. In particular, interaction is a joint process of user with the other components (also of the other components with one another), and judgment is a process which can only be performed by the user. Furthermore, although modification is something that can be done by the other components of the system with reference to modifying query or text representation, modification of understanding of the information problem is something that can realistically be done only by the user. Thus, the idea of the 'intelligent intermediary' as being the basis of intelligent IR, although perhaps necessary, is not sufficient to characterize the complete intelligent IR system. Similarly, the idea of good IR as being effective IR fails if all the intelligence is concentrated in only the built system, since it thereby excludes the most significant aspect of effectiveness, the user's judgment of the comparison performance.

These arguments lead us to the position that intelligence in intelligent IR perhaps resides not only in that which is built to support the user, but also in the user her/himself, and in particular in proper assignment of roles and responsibilities to all of the various components of the IR system. Clearly, some of the processes we have mentioned must be carried out by the components in concert, and some knowledge on which they are based (for instance, knowledge of database structure and contents) might be privileged to only one of the built components. Thus, there is some sense in which we might say that doing these processes (or applying such knowledge) well is important to intelligent IR. However, this view does not incorporate the undoubted intelligence of the user, in particular in respect of the judgment, modification and interaction processes. Thus, we might say that a role of the user is to judge or evaluate the texts with respect to the information problem; a role of the intermediary is to offer to the user interesting texts with which to interact; and a role of the knowledge resource is to be organized in such a way as to promote effective interaction. These are only examples of the roles and responsibilities of the various actors (perhaps a better term than components) in the IR system, but they suggest how we might go about construing truly intelligent IR. That is, by first considering the essential nature of the IR situation, and then designing the IR system to promote and enhance the activities of all of the actors in the system with respect to that essential nature.

4. Information retrieval systems as support for interaction with information.

In Belkin 93, I suggested a view of IR as information seeking behavior, a kind of interaction with text. Here I would like to propose an extension of that view, which allows us to consider

IR systems as systems for supporting people's interactions with information. This position can, I think, lead us to a better understanding of IR in general, to some interesting ideas for the design of IR systems, and perhaps even to a new way to construe intelligent IR. The key to this proposal is that it attempts explicitly to make people's interactions with information the central process of IR, with the other processes and components being seen as providing methods for the appropriate support of such interaction.

The basis for this view of IR is the observation that people engage in a wide variety of information seeking behaviors, and more generally, interactions with information, both in different information-seeking episodes, and within the course of a single information seeking episode. This observation has led, for instance, to a potential classification of information seeking behaviors or strategies (ISSs), based upon observable characteristics of human behavior in interaction with information. It has also led to the idea that for each such ISS, there might be a prototypical or 'best' way to accomplish it, within the constraints of an IR system. In Belkin et al. 93 and Belkin et al. 95, these ideas have been used, respectively, as the basis for the design of an IR system interface which allows easy movement from one kind of ISS to another; and, for designing a dialogue-based IR system which actively supports different kinds of interaction for different kinds of ISS, again with easy movement from one to another. Here, I will not describe the details of these proposals and systems, and ask the reader to accept that at least some reasonable work of this type has been done.

Although this work has gone some way toward understanding and incorporating different kinds of information seeking within IR system design, it is still somewhat hampered by the lack of some way both to understand when some specific kind of support was needed, and to place interaction, rather than the other processes of IR, and the user, rather than the other components, at the center of the IR system. Below, I sketch an outline of what seems to be a way to accomplish these goals.

If we consider the IR system as a system whose goal is effective support of a person's interactions with information (in particular, those associated with information-seeking behavior), then it is necessary for us to understand at least the following:

what are the *kinds* of interactions in which people engage?

what *situations* or *contexts* or *goals* lead to specific kinds of interactions?

how does the nature of the *information objects* interacted with affect the nature of the interaction itself?

In addition, we might reasonably expect that we should know something about the sequential nature of an information seeking episode, in particular what might lead to change from one kind of interaction to another. And finally, we will need to know whether there are some different ways to support optimally different types of interactions. Figure 1 presents a general model of IR as support for information interaction which attempts to take all these issues into account.

We read figure 1 as follows. At any point in time in an information seeking episode, a person will be engaged in some specific kind of interaction with some specific kind of information object. The kind of interaction, and perhaps the kind of information object interacted with, will be dependent upon that person's goals, problem, intentions, situation, etc. at that time, and on the course of the interaction to that point. Such information interaction is supported by a variety of processes, or actors other than the user in the IR system. Such processes include, for instance: representation, comparison, presentation, navigation, visualization, and so on. Each such process can be instantiated by one of several different techniques; we hypothesize that, for any particular kind of interaction, there will be some optimum combination of techniques from the various processes, for effective support of that kind of interaction.

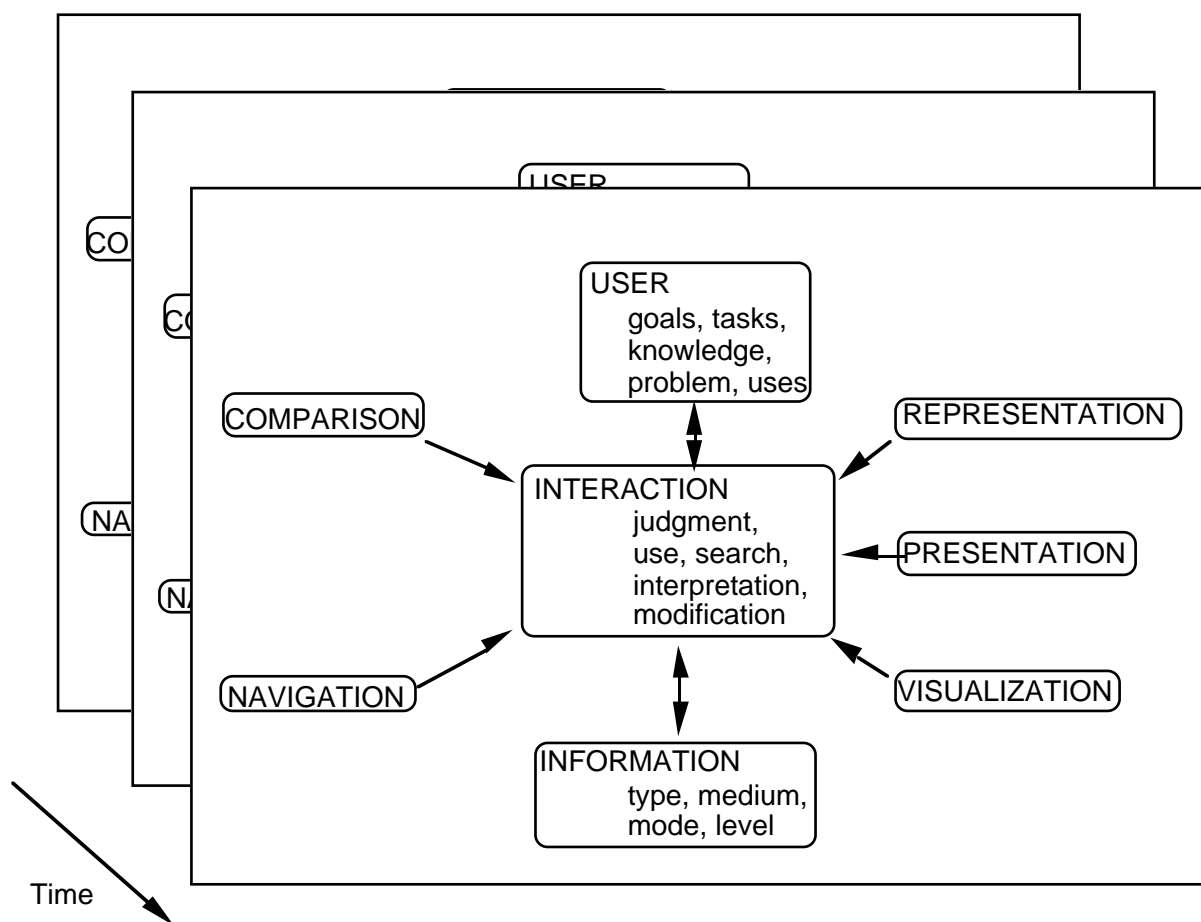


Figure 1. Information retrieval as support for information interaction

An information seeking episode consists of a series of kinds of interactions (slices in time), structured according to some plan associated with the person's overall goals, problem, experience, according to the person's specific goals, etc. at any one time, and according to what has happened during the course of the interaction. For instance, a person with the general goal of learning about a new topic might initiate the IR system by interacting with some meta-data resource, in order to learn about the contents of the available database(s). The person might then put a specific query to a database, in order to learn whether there are documents in it which might be relevant to her problem. Having perused some documents found in this way, and having judged them all to be non-relevant, the person might begin to explore, perhaps through a thesaurus, other ways that the concepts in which she is interested. Finding one such concept that seems likely, she looks at a document which is indexed by that concept. Liking what she sees, she looks through some other documents that seem closely related to that one. On the basis of some relevance judgments on these documents, a query is generated to search again in the database. The documents which are retrieved by this comparison process are presented to the person as a set of classes of related documents. The user, judging one of the classes to be quite interesting, asks for a summary of those documents. A summary is presented to the user, which gives her enough information so that she can do the task which lead to her goal of learning about a new topic, and the episode is terminated.

This little scenario is meant as a demonstration of the dynamic and changing nature of interaction with information during the course of an information seeking episode. At each point, a different kind of information seeking behavior is taking place, conditioned by both the original goal, knowledge, problem, and by what has happened to that point. It also demonstrates how the different kinds of interactions are best supported by different combinations of differ-

ent techniques from each of the IR support processes. And, finally, it suggests, I think, how one might construe intelligent IR.

In such a scenario, the user plays a central role, guiding the system, making evaluative judgments, deciding about what to do and when to stop. The other processes contribute by understanding something about what is likely to help the user in supporting the interactions in which that person is engaged, in knowing something about what the likely course of the interaction as a whole might be, and in using their knowledge about the resources at their disposal to inform the user about the system and its contents so that the user can interact effectively. Thus, using this model, intelligent IR turns out to be IR in which intelligence is explicitly distributed throughout the system, all of the actors contributing according to their specific roles and knowledge to support the user's effective interaction with information.

References

Bates 90: M. Bates. Where should the person stop and the information search interface start? *Information Processing and Management*, 26(5): 575-59, 1990.

Belkin 84: N.J. Belkin. Cognitive models and information transfer. *Social Science Information Studies*, 4(2-3): 111-129, 1984.

Belkin 93: N.J. Belkin. Interaction with texts: Information retrieval as information seeking behavior. In: *Information Retrieval '93: von der Modellierung zur Anwendung*. Regensburg, 1993, pp. 55-66. Konstanz, Universitätsverlag Konstanz, 1993.

Belkin and Vickery 85: N.J. Belkin and A. Vickery. *Interaction in information systems*. London, The British Library, 1985.

Belkin, et al. 83: N.J. Belkin, Th. Seeger and G. Wersig. Distributed expert problem treatment as a means for information system analysis and design. *Journal of Information Science*, 5: 153-167, 1983.

Belkin et al. 87: N.J. Belkin, H.M. Brooks and P.J. Daniels. Knowledge acquisition using discourse analysis. *International Journal of Man-Machine Studies*, 27: 127-144, 1987.

Belkin et al. 93: N.J. Belkin, P.G. Marchetti and C. Cool. BRAQUE: Design of an interface to support user interaction in information retrieval. *Information Processing and Management*, 29(3): 325-344, 1993.

Belkin, et al. 95: N.J. Belkin, C. Cool, A. Stein and U. Thiel. Cases, scripts and information seeking strategies: On the design of interactive information retrieval systems. *Expert Systems with Applications*, 9(3): 379-395, 1995.

Croft 87: W.B. Croft. Approaches to intelligent information retrieval. *Information Processing and Management*, 23 (4): 249-254, 1987.

Ingwersen 92: P. Ingwersen. *Information retrieval interaction*. London, Taylor Graham, 1992.

Ingwersen and Wormell 86: P. Ingwersen and I. Wormell. Improved subject access, browsing and scanning mechanisms in modern online IR. In: F. Rabitti, ed. *SIGIR '86. Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*, June 1986, Pisa, pp. 68-76. Pisa, ACM, 1986.

Maes 94: P. Maes. Agents that Reduce Work and Information Overload. *Communications of the ACM*, 37(7): 31-40, 1994.