From E-Relevance to W-Relevance

Paolo Coppola, Vincenzo Della Mea, Luca Di Gaspero, Stefano Mizzaro
Department of Mathematics and Computer Science
University of Udine
http://www.dimi.uniud.it/~{coppola,dellamea,digasper,mizzaro}
{coppola,dellamea,digasper,mizzaro}@dimi.uniud.it

Abstract. We discuss how the wireless-mobile revolution will change the notion of relevance in information retrieval. We distinguish between classical relevance (e-relevance) and relevance for wireless/mobile information retrieval (w-relevance). Starting from a four-dimensional model of e-relevance previously developed by one of us, we discuss how, in an ubiquitous computing environment, much more information will be available, and how it is therefore likely that w-relevance will be more important than e-relevance to survive information overload. The similarities and differences between e-relevance and w-relevance are described, and we show that there are more differences than one might think at first. We specifically analyze the role that beyond-topical criteria have in the w-relevance case, and we show some examples to clarify and support our position.

Keywords: Relevance, mobile, wireless, ubiquitous, e-relevance, w-relevance, beyond-topical criteria, topic, task, context.

1. Introduction

We can hardly imagine what information overload is. If this sentence seems surprising to you, just stop one minute and think how our world will probably be in ten years or so. As soon as mobile wireless devices will ubiquitously enter our lives, the nowadays complaints about having access to too much information will be seen with a small ironic grin and perhaps some nostalgia. We are not speaking only of palm top devices, cellular phones, laptop computers, pagers, MP3 players, and similar already commonly used device; we are thinking also of networked digital cameras and video-cameras, thermometers, traffic lights, GPSs for cars and – why not – bikes, skates, pogosticks, and even walking people, game stations, and so on. Thousands of interconnected information processing devices will be available to each of us anytime anywhere. Each mobile device will sense its environment to gather information from the physical world and make it available to its user (or users). Each device will also exchange information with other (mobile and non-mobile) devices, mainly by means of some wireless communication network. Probably, users will (continue to) directly exchange information among them. Also, devices will probably change the physical environment, much more than nowadays static and non-ubiquitous desktop machines. A similar view is expressed, for instance, in [Lettieri & Srivastava 1999].

All the mobile devices can be seen, from the user point of view, as information access tools: they will filter incoming information and retrieve available information, trying to present to the user all and only the relevant information. Of course, the user will be interested in accessing information that is not only relevant in the strict sense, but also of a high quality, timely, serendipitous, of the appropriate grain size, perhaps rare, and so on. Since there is not an agreement about which of these features are relevance features, we will use the term "relevance" in a very general way, denoting with relevant information the information that the user wants.

But what is relevance in the new mobile/wireless/ubiquitous scenario? This paper is a first and preliminary attempt of answering this question. In Section 2 we will briefly overview the research about the concept of relevance in classical non-mobile *Information Retrieval* (henceforth IR). We name relevance in classical IR *e-relevance* (for electronic relevance, but this is not the only reason, as we will explain in Section 5). In Section 3 we will re-analyze the relevance concept in the mobile case. In turn, we name this relevance *w-relevance* (for wireless relevance, but, again, see Section 5). We show that, from an intuitive point of view: (i) w-relevance is an extension of e-relevance; (ii) w-relevance is much different from e-relevance than one might think at first; and (iii) beyond-topical criteria, one aspect of e-relevance that has recently received a lot of attention in non-mobile IR, are both much more emphasized and much more important in the mobile case. In Section 4 we propose some simple examples and scenarios to support our position. Section 5 concludes the paper and fully explains its title.

2. E-Relevance: The Non-Mobile Information Retrieval Case

Relevance (e-relevance) is a subject that has been intensely studied for years in the IR field, and it is still a hot topic today. We will not review in detail the field, since some well known surveys are already available [Saracevic 1975] [Schamber, Eisenberg & Nilan 1990] [Schamber 1994] [Saracevic 1996] [Mizzaro 1997].

Classical information retrieval equates e-relevance with topicality: the query submitted to an IR system specifies the topic(s) that a relevant document has to deal with. For example, if a university professor is looking for documents to prepare her next lesson for this afternoon, she needs of course documents that deal with the matter that she is going to explain to her students.

But she also wants those documents as soon as possible (if a document arrives after the lesson, it is useless), at the right complexity level (if a document is too difficult, students will not understand it), and so on. And these features go beyond the topic: they are completely independent of it.

Therefore, the topical view is short-sighted. Indeed, we have now a large amount of research that demonstrates how topic is only one of the criteria that users use when judging the e-relevance of the retrieved documents. For a review of this line of research, that started in the 60es and has received a lot of attention, especially at Syracuse University, in the 80es and 90es, see [Mizzaro 1997] [Barry & Schamber 1998]. Since the criteria, elicited from users or found by experts, tend to constitute a stable set (i.e., very few new criteria are found in the most recent studies), it is likely that we have an almost correct and complete list of relevance criteria.

Actually, the exploitation of beyond-topical criteria is not the only way to get closer to the "real" relevance, i.e., the relevance the user is interested in. A more general approach that takes into account this aspect has been proposed by one of us some years ago [Mizzaro 1998] [Gabrielli & Mizzaro 1999]: the various kinds of relevance are classified in a four-dimensional space, distinguishing among them on the basis of a precise classification. The four dimensions are:

- *Information resources*, containing document, surrogate, and the information that the user receives when reading a document.
- *Representation of the user problem*, containing the real information need, the perceived information need, the request (or expressed information need), and the query (or formalized information need).
- *Time*, containing the time instants from the arising of the user's need to its satisfaction.

- *Components*, containing topic, task (what the user has to do with the retrieved information), and context (everything beyond topic and task as, for example, what the user already knows about the topic being sought, or the time that the user has to complete the search).

These four dimensions allow to distinguish among the various kinds of relevance, and to speak, for instance, of: the relevance of a document to the query at query expression time for what concerns the topic component (the classical relevance used in IR); the relevance of the information received to the real information need at the time of final need satisfaction for what concerns topic, task, and context (the relevance the user is interested in); and so on. This classification can be used in the implementation and evaluation of IR systems.

This topic/task/context distinction has been used in some respect. [Reid 2000] proposed an evaluation methodology that uses the task as the starting point for building a test collection. The development of IR systems dealing with beyond-topical e-relevance has been rather slow, however some examples now exist. Researchers at MIT recently developed an IR system that, in some way, goes beyond topical criteria [Liu, Lieberman, & Selker 2002]. This system, named GOOSE (GOal Oriented Search Engine), allows the user to choose among a list of tasks (called "goals" by GOOSE authors), and uses a large common sense knowledge base to exploit the task specification for building a better query. In such a way, Liu and colleagues implemented, perhaps without explicitly noting it, an IR system that tries to work taking into account beyond-topical factors of relevance, as suggested in [Mizzaro 1998].

One can also assume that, although each search and each information need concern a different topic, there are indeed some beyond-topical components of user's needs that are more stable, i.e., the context in which the consecutive search sessions by one user take place [Göker 1997] [Göker 1999]. Some first experiments show that, for a given user, contexts are indeed more stable than topics, and may be used to improve the ranking of documents retrieved after a query, but the usefulness of this approach is still under investigation.

Another approach for including beyond topical criteria in an IR system is to build an *IR assistant*, namely a system that, during information seeking, observes user behavior and gives suggestions aimed at improving the effectiveness of the search and of the searcher [Brajnik, Mizzaro & Tasso 1996] [Brajnik et al. 2002] [Mizzaro & Tasso 2002]. Some of the suggestions might be of a topical nature (e.g., to add some terms to the query to better represent the topic being sought for), but also non-topical suggestions can be provided, like suggesting a paper related in some way to those judged as relevant so far (e.g., the PhD thesis by, or a short biography of, the author of a paper judged as relevant, or a references list, and so on). This line of research has still to be proven effective, but initial laboratory experiments show positive results. Also "just-in-time information retrieval agents" [Rhodes & Maes 2000] build their queries with beyond-topical components (mainly context).

Even if the existence of beyond-topical criteria for e-relevance is not in discussion, what seems not yet recognized, or assessed, is the actual importance of these criteria in real-life IR. In the next section we discuss, on the basis of the classification in [Mizzaro 1998], how and why the w-relevance scenario is different.

3. W-Relevance: The Mobile Information Retrieval Case

One might simply repeat the above analysis in the w-relevance case, and thus just state that there are various kinds of w-relevance and there are some beyond-topical components of w-relevance that should not be overlooked. However, we believe that there are important differences between e-relevance and w-relevance. The beyond-topical criteria in the mobile IR case become more critical: they are different from, and have a higher importance than, those in non-mobile IR. Therefore, using only topicality is an abstraction, a simplification, that works in a perhaps satisfying (even far from perfect) way in the e-relevance case but, as soon as the real world comes into play, as it happens in the w-relevance scenario, the shortcomings of this approach are manifest (examples will be shown in Section 4). Also, there are more kinds of w-relevance than kinds of e-relevance. As we will discuss in the following, the main reason for these differences is that in the e-relevance case we can comfortably seat inside the "information world", whereas in the w-relevance case we have to move into the "real/physical world".

All the e-relevance models proposed in past years need to be modified to become adequate models of w-relevance. In this section we revise and extend the model proposed in [Mizzaro 1998], in each of the four above mentioned dimensions.

3.1. Information resources

In the non-mobile case, the user of an IR system is usually interested in retrieving information; a typical user is a scholar that needs documents on a new topic, to study them, to write a paper or book, and so on. This is obtained by retrieving a number of information sources, from which the relevant information can be extracted. In the mobile IR scenario, it is often the case that the user is interested not just in information, but in obtaining some (possibly material) *thing*, only partially described by information: besides surrogate, document, and information, the information resources dimension should therefore contain also thing. In other terms, often the retrieved information and, in general, the database are instrumental, they are means to reach the end of having some thing, not the end itself. Besides the relevance of retrieved information, we also have the relevance of the retrieved thing: the user will not evaluate the information sources, but the described physical object, that in the meantime might change or disappear even without an immediate reflection on the information source content. This brings up the issue of consistency between the database and the real world.

3.2. Representation of the user problem

Since in the mobile IR scenario, besides the real information need (that is in turn beyond the information need perceived by the user), it is often the case that the user is interested in some *thing*, we can say that the user usually has a *thing need* (that should then be added to the second dimension, namely the representations of the user problem). Therefore, if we look at the first two dimensions, we can say that from the relevance of the retrieved information to the information need, we have moved to the relevance of the retrieved thing to the thing need. Using Bateson's [Bateson 1979] terminology, w-relevance deals more with *Pleroma* (the physical world), whereas e-relevance deals mainly with *Creatura* (the informational world): in w-relevance we have a much stronger coupling with the real, physical world. If one photocopies an article in a library (e-relevance scenario), you can anyway read the article later. If someone buys the last item of your favorites blue jeans just after your query to a "blue jeans database", you cannot have them anymore (w-relevance scenario). In the former case you are interested in information, whereas in the latter you are interested in a thing.

3.3. Time

Another dimension of relevance that increases its importance in the w-relevance case is *time*, in two senses. First, often the user needs "quick and dirty" information: things change faster, replication is more difficult. Second, in the real world, since time is irreversible, if something is lost it is lost. In the Creatura one can often rely on backups, copies, and replication; in the real world, "carpe diem". This is perhaps the deep motivation behind the often stated claim that users of mobile devices are more interested in precision than in recall, usually justified, in a perhaps too simplistic way, by the small display area on mobile devices: having a full list of the relevant items can be useless if the list is so long that the time required for examining it is longer than the lifetime of relevant items.

Another aspect of w-relevance related to the strong coupling with the real-world and to time is the change rate of the database: since the real world changes quickly and continuously, the database has to quickly change accordingly to stay up-to-date.

The intuitive importance of time is also confirmed by a survey made last year in Singapore among users of PIRO, a commercial system developed by C5solutions [C5solutions]. PIRO presents to mobile users using WAP phones the directory listings of commercial retail relevant to user's current need. Eight users filled in a 29 questions questionnaire having the purpose to rate the importance of various relevance criteria for presenting commercial applications on a mobile device. Of course the small sample size does not allow any certain inference, but it is worth to note that two out of the three most higher rated criteria are "information is current (up to date)" and "information is about a sale or promotion or money saving opportunity", both of which concern time features.

3.4. Components

"Context" is a hot word in the mobile/wireless scene, but with a different meaning from that used above [Chen & Kotz 2000] [Cheverst et al. 1999] [Cheverst et al. 2001] [Pascoe 1998]: context usually refers to the current environment, the situation, that the user of a mobile device is experiencing while using it. Let us see some examples of this usage of the term.

Location is one of the most mentioned aspects of context [Ward, Jones & Hopper 1997]: from the user position (derived by means of GPS, or triangulation in a Bluetooth or Wi-Fi network) other information can be inferred and exploited in various ways, for example to increase the relevance of the information accessed by the user, or to improve the interaction with the user.

Of course, location is important, but there is more to context than location [Schmidt, Beigl & Gellersen 1999]. Location itself it is not the only information we can get from the spatial position of the user. Indeed, this feature is only as a "static" one, whereas several additional information can be inferred from the dynamic evolution of locations. For example, if the user looks for traffic information, and she is moving along a road, it is very likely that she wants information about the road she is currently on, rather than the whole national traffic news. Therefore, user's track, i.e., the temporal sequence of locations traveled by a user, is another aspect of context. Let us notice that context can also be predicted; for example, a full track can be inferred if the user's scheduler reveals that she has an appointment in half an hour at a certain place. Also the traveling speed that the user has while following a certain track is an important parameter: a slowly walking user can be presented more information than a running one [Wahlster 2002].

Other common examples of context aspects are: the noise level in the environment (that can and should affect the volume of a mobile phone); the light level in the environment (affecting the display illumination); the orientation of the device (affecting the orientation of the displayed information) [Schmidt, Beigl & Gellersen 1999].

Therefore, "context" has a different meaning in w-relevance: in e-relevance, context concerned what was in user's mind only; whereas in w-relevance, context is also about the real world. We will distinguish between these two meaning by using *e-context* and *w-context* terms. W-context is more general. Moreover, w-context is much more dynamic than e-context, and it is more likely to change during the information seeking activity. It is also worth noting that the above mentioned Göker's preliminary results on context stability might fail in the mobile environment. Finally, e-context has to be provided manually to an e-relevance IR system, whereas w-context can and should be derived autonomously.

4. W-Relevance Scenarios

In this section we show some realistic scenarios that support the above discussion. We are aware of the importance of privacy issues, but we do not take them into account in this paper.

4.1. Catching the train

Let us consider a query for finding a train from Udine, Italy to Milano, Italy, made to the Italian national railways web site (http:// www.trenitalia.com). At present, you fill a form with (at least) departure and arrival cities, and starting date and time. If the two cities have more than one railway station, you are also asked to select the specific stations you want. Then, you receive a list of trains since that hour, and if you want details on one of them, you have to click on its link, then go to the price link and further compile a form where you specify how many tickets, in which class, etc.

What if your are reaching the station in a hurry on a taxi, just in time for a train? Your need is not to be informed about trains from Udine to Milano: you just need to go where your are used to go, Milano Piazza Garibaldi railway station, in second class as usual, by the first train available. Other options include: you are not alone, but with your husband/wife (your PDA is sensing him or her around you); the ticket can be bought automatically (using the available details of your credit card); and the first train retrieved could be not useful because there are not two free seats in second class. All these data can be derived from your own w-context and an up-to-date (with respect to the real world) database..

4.2. Driving to SIGIR 2012

Let us imagine that you are driving your car, rented at Helsinki airport, towards Tampere to attend SIGIR 2012. Your car is of course equipped with a GPS and a driving assistant, giving you directions about the route to follow. In this situation, the information that the car 100 meters in front of you is driving to Tampere too is very relevant for you, and should be immediately notified so that you might follow that car without worrying about road directions. Moreover, if the driver in the other car is a good friend of yours, you should be notified about that too, since you might want to contact her (with an SMS?) for sharing the trip or just having a coffee together.

In this scenario, the topic is straightforward, being the destination of your trip (Tampere); the task is given mainly by driving in a convenient way, with perhaps some subtasks given by sharing the trip for economy, avoiding pollution, just chatting. The latter case is even more interesting if the driver is a friend, but this is not expressed neither in topic nor in task, but most likely in the e-context – your address book, your last phone calls, etc. You might go on, and think of the situation if your good friend is not a good driver at all, or if her car is a very old one (and these are w-context aspects).

4.3. E-commerce application

Now you are a trendy boy/girl, shopping around in a commercial center, and willing to buy some fashionable trousers (“Gasoline” brand), and a newly available “Mosquito” shirt. You do not want to spend too much money, so you ask your PDA to look for those dresses at a good price. Some different outcomes might be considered: (i) your are using a traditional e-relevance based IR system, thus you will receive a list of offers on eBay, followed by some online shop catalogue showing very good prices for the same dresses; unfortunately, you are around for shopping, you want to have your dresses now and not to wait for their postal delivery; (ii) your PDA queries a w-relevance based IR system which, on the basis of your location, track, and walking speed, is able to infer that you probably want some specific place where to buy such dresses, possibly close to you.

So you will receive three shop addresses: the closest one is not the first because it is slightly more expensive than another one, which in turn is sufficiently near to be reached before closing time. A third shop is listed with good prices but with just one shirt of your size (as recorded in your personal profile, or communicated after request by the radio/infrared label applied on the shirt you are actually wearing); you have to run before someone other buys it, or perhaps you can book it by means of an electronic message. Again, the topic is Gasoline trousers, Mosquito shirt, good price, but the real task is to actually buy them, not to know where to buy them. Real things get sold out, usually on a first-came first-served basis: what is true now (e.g., availability of my size) could be false in some minutes, thus time matters too.

4.4. A museum application

Just in front of a beautiful Van Gogh picture, you want to have some more information to understand why he is painted with a bandaged ear. The wireless service available there of course does not need to show you the picture itself: you just need textual background information, as in you current w-context there is the real availability of the object which generates the topic for the query. In the wireless-enabled environment, your PDA should handshake with a radio/infrared label applied near the picture, so that, in addition to exactly know your position, it is also able to automatically generate part of the query.

5. Conclusions and Future Work

In this paper we have shown how the mobile/wireless/ubiquitous revolution is likely to bring big changes into the IR field, and how even a very foundational concept as relevance needs to be re-analyzed and defined.

The title of this paper reflects this claim: the relevance in classical electronic environments (non-mobile ones) that we have named *e-relevance*, is actually an *irrelevance*, because many features of it are neglected, or at least not given the importance that they should have in the general case. In the mobile IR case, this generality is more easy to notice: *w-relevance* does not only mean wireless relevance, but also *double-relevance*, *world-relevance* (since the physical world is much more involved) and *double-user-relevance*, since it is a notion of relevance that is much more close to what the users want and need.

The current model for information retrieval – with one user and one system – is also challenged in a peer-to-peer wirelessly connected environment, where ultra-mobile devices are available for providing information to each other. Information will be available from many devices through many channels, either phone-like (WAP, GPRS, UMTS) or local networking (Wireless LAN, Bluetooth, IrDA). Such devices may in turn provide also w-context information, i.e., location (in a broad sense, not only geographical coordinates), track, temperature, etc. In such a complex peer-to-peer scenario, it is likely that a single query made by my device could be answered by more than one system, and that each system could be engaged in a sort of "reverse relevance", asking to itself something like "Am I able to answer to such a query?", which in turn could be translated as "Is such a query relevant to my database?". As device answers may have a cost for the user, it is also likely that the query should involve budgetary considerations. Useful hints about how to deal with such a kind of interactions may come from the multi-agent paradigm research area [Wooldridge 2002].

Moreover, another problem needs to be mentioned. On the one side, it seems reasonable that as soon as some information is available and potentially relevant in the future, it should be stored locally on one's own device to be accessible later. This is even more reasonable if one takes into account that wireless devices are not always connected to the network, and that they use different network connections, with different transfer rates, reliability, privacy, and cost. However, on the other side, small devices are more resource constrained: low computational and storage power, low energy availability, low bandwidth also in the interaction with the user.

All these features would suggest that the local storage of information is not always the best choice. Also, the locally stored data can quickly become outdated because of the quick change rate of the database, therefore rising inconsistency problems. This issues need further investigation.

In the future we plan to work on the relevance four dimensional model in order to make it more accurate and formalized. We also intend to use the revised model in the implementation and evaluation of mobile IR systems.

Acknowledgements

We would like to thank Roberto Ranon for interesting discussions and useful comments on an earlier draft of this paper.

References

[Barry & Schamber 98] Barry, C. L., Schamber, L. Users' Criteria for Relevance Evaluation: A Cross-Situational Comparison, *Information Processing and Management* 34(2/3), 1998, 219-236.

[Bateson 1979] Bateson, G. *Mind and Nature – A Necessary Unit*. E. P. Dutton.

[Brajnik et al. 2002] Brajnik, G., Mizzaro, S., Tasso, C., Venuti, F. Strategic help in user interfaces for information retrieval, *Journal of the American Society for Information Science and Technology*, 53(5), 2002, 343-358.

[Brajnik, Mizzaro & Tasso 1996] Brajnik, G., Mizzaro, S., Tasso, C. Evaluating User Interfaces to Information Retrieval Systems: a Case Study on User Support, *Proceedings of the 19th Annual International ACM SIGIR Conference*, Zurich, CH, 1996, 128-136.

[C5solutions] Insuring an Industry, Investing in People. *Career Times* 2001/11/09. http://www.careertimes.com.hk/english/employers/mgmt/features/20011109.asp

[Chen, C. & Kotz 2000] Chen, C., Kotz, D. *A Survey of Context-Aware Mobile Computing Research*. Dept. of Computer Science, Dartmouth College Technical Report TR2000-381, November 2000, ftp://ftp.cs.dartmouth.edu/TR/TR2000-381.ps.Z.

[Cheverst et al. 1999] Cheverst, K., Davies, N., Mitchell, K., Friday, A. The role of connectivity in supporting context-sensitive applications, *Handheld and ubiquitous computing, Proceedings, LNCS, 1707*, 1999, 193-207.

[Cheverst et al. 2001] Cheverst, K., Smith, G., Mitchell, K., Friday, A., Davies, N. The role of shared context in supporting cooperation between city visitors, *Computers & Graphics*, 25(4), 2001, 555-562.

[Gabrielli & Mizzaro 1999] Gabrielli, S., Mizzaro, S. Negotiating a multidimensional framework for relevance space. In S. W. Draper, M. D. Dunlop, I. Ruthven, and C. J. van Rijsbergen, editors, *Mira 99: Evaluating interactive information retrieval - Proceedings of MIRA 1999 Conference*, eWiC - electronic Workshops in Computing, pages 1-15, 1999.

[Göker 1997] Göker, A. Context Learning in Okapi, *Journal of Documentation*, 53(1), 1997, 80-83.

[Göker 1999] Göker, A. Capturing Information Need by Learning User Context, *Sixteenth International Joint Conference in Artificial Intelligence: Learning About Users Workshop*. 1999, 21-27.

[Lettieri & Srivastava 1999] Lettieri P., Srivastava, M.B. Advances in wireless terminals, *IEEE Personal Communications*, 6 (1), 1999, 6-19.

[Liu, Lieberman & Selker 2002] Liu, H., Lieberman, H., Selker, T. GOOSE: A Goal-Oriented Search Engine with Commonsense. In Paul De Bra and Peter Brusilovsky editors, *Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Malaga, 29-31 May 2002, 253-263.

[Mizzaro 1997] Mizzaro, S. Relevance: The whole history. *Journal of the American Society for Information Science*, 48(9), 1997, 810-832.

[Mizzaro 1998] Mizzaro, S. How many relevances in information retrieval? *Interacting With Computers*, 10(3), 1998, 305-322.

[Mizzaro & Tasso 2002] Mizzaro, S., Tasso, C. Ephemeral and Persistent Personalization in Adaptive Information Access to Scholarly Publications on the Web. In Paul De Bra and Peter Brusilovsky editors, *Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Malaga, 29-31 May 2002, 306-316.

[Pascoe 1998] Pascoe, J. Adding generic contextual capabilities to wearable computers, in *Proceedings of 2nd International Symposium on Wearable Computers*, October 1998, pp. 92-99.

[Reid 2000] Reid, J. A Task-Oriented Non-Interactive Evaluation Methodology for Information Retrieval Systems, *Information Retrieval*, 2(1) , 2000. 115-129.

[Rhodes & Maes 2000] Rhodes, B. J., Maes, P. Just-in-time information retrieval agents, *IBM Systems Journal*, 39(3&4), 2000, 685-704.

[Saracevic 1975] Saracevic, T. Relevance: A review of and a framework for the thinking on the notion in information science. *Journal of the American Society for Information Science, 26,* 1975, 321–343.

[Saracevic 1996] Saracevic, T. Relevance Reconsidered ’96. In P. Ingwersen and N. O. Pors, editors, Information Science: Integration in Perspective - Proceedings of CoLIS2, pages 201-218, Copenhagen, Denmark, October 1996. The Royal School of Librarianship.

[Schamber 1994] Schamber, L. Relevance and information behavior. *Annual Review of Information Science and Technology,* 29*,* 1994, 3-48.

[Schamber, Eisenberg & Nilan 1990] Schamber, L., Eisenberg, M. B., Nilan, M. S. A re-examination of relevance: Toward a dynamic, situational definition. *Information Processing & Management,* 26(6), 1990, 755–776.

[Schmidt, Beigl & Gellersen 1999] Schmidt, A., Beigl, M., Gellersen, H.-W. There is more to context than location, *Computer & Graphics Journal*, 23(6), 1999, 893-902.

[Wahlster 2002] Wahlster, W. Resource-Adaptive Interfaces to Hybrid Navigation Systems (Keynote Talk), In Paul De Bra and Peter Brusilovsky editors, *Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Malaga, 29-31 May 2002, 12-13.

[Ward, Jones & Hopper 1997] Ward, A., Jones, A., Hopper, A. A New Location Technique for the Active Office. *IEEE Personnel Communications*, 4(5), 1997, 42-47.

[Wooldridge 2002] Wooldridge, M. *An Introduction to Multiagent Systems*. John Wiley & Sons, 2002.