ELECTRONIC WORKSHOPS IN COMPUTING

Series edited by Professor C.J. van Rijsbergen

Ian Ruthven (Ed)

Miro'95

Proceedings of the Final Workshop on Multimedia Information Retrieval
(Miro '95)
Glasgow, Scotland
18-20 September 1995

Paper:

Supporting Information Needs by Ostensive Definition in an Adaptive Information Space

Iain Campbell

Published in collaboration with the
British Computer Society

©Copyright in this paper belongs to the author(s)

Supporting Information Needs by Ostensive Definition in an Adaptive Information Space

Iain Campbell

Computing Science
University of Glasgow
Glasgow G12 8QQ

iain@dcs.gla.ac.uk

Abstract

In this report I describe those characteristics of information needs that I consider to be the most important for their effective modelling in Information Retrieval (IR). Using these characteristics as terms of reference, I highlight inadequacies of current IR approaches. I propose a new approach that is not only a hybrid of positive aspects of the existing approaches, but that also embodies a significant and novel addition - the ostensive definition of developing information needs. The new approach displays interesting properties - eg the contextual interpretation of information items. I describe these properties and explain why I feel it makes a more appropriate model for the provision of an environment in which to perform searching tasks. I believe that it will facilitate less formal, more spontaneous access to stored information, and that it will reduce the difference in degrees of effectiveness attained by users of varied experience. It has important media, language, and domain independence characteristics that make it particularly suitable for the next generation of information systems. The approach, centred around a graphical interaction environment, brings with it new, possibly unique, problems in visualisation - challenging one of the accepted assumptions of information spaces.

1 Information needs

When first setting out to track down information, be it an answer to a specific question, or general information on a topic of interest, one often perceives of an 'information need'. Such information needs are nebulous combinations of ideas such as what the target information might look like, where it might be found, or how one might go about tracking it down - with the words *look*, *where*, and *how* used in a most general sense. It is difficult, both for an observer and for the person involved, to determine the information that constitutes those ideas, and further, to determine the various proportions of each that make up the currently perceived information need.

Attempting to describe an information need is a difficult enough task to perform effectively to another intelligent human being. Upon casual analysis, then, trying to do the same to a simplistic, uninformed, inexperienced computer system strikes one as plainly ridiculous. Putting the glib comment regarding computer systems' capabilities to one side, that thought hides within it two ideas that I regard as key to understanding the nature of information seeking. The first is that of the whole idea of describing an information need, and the second is our apparent willingness to make such descriptions to systems known to be so seriously flawed.

1.1 Describing information needs

Before talking further about an information need, perhaps it would be better to consider an information *lack*. I believe it is the perception of a lack of information that provokes one to then develop a need for it. It is a simple fact that information one does not currently have is information that one cannot describe. A full description of something cannot be *made* by anything, it can only *be* the item itself. The existence of something, be it conceptual or physical, *is* the only full description of that thing. We can attempt to describe it in many ways, for example: Physically - in terms of light reflected from it, sound emitted from it etc; Compositionally - in terms of its constituent parts; Functionally - in terms of what it can or cannot do; Procedurally - in terms of how it does things. Such descriptions can only be partial descriptions - ie reliant upon some restricted view of the world, and as such, laden with assumptions and simplifications.

All physical objects can be thought of as information - eg a particular arrangement of elementary particles and waves (be that arrangement absolute or probabilistic). Concepts can also be thought of as particular arrangements of other, perhaps elementary, concepts. To these ideas of arrangement, one can add the idea that a particular

arrangement may change over time and still be regarded as the same thing. All these arrangements, and arrangements of arrangements are information. Everything can be regarded as information, and hence information can be thought of as a *thing* in the purest and most general sense. Finally, recognising that a particular arrangement can only *be* something if an observer regards it as such, the 'nature' of a thing can be recognised as a totally subjective concept and impossible to tie down.

From the above, it can be seen that information that is perceived to be lacking can neither be known nor described. Only the fact that information is missing can be recognised. We often make estimates as to the size of this lacking, but as no clear concept of a quantitative measure for information has been developed or distilled from observation, such estimates have little meaning. In fact, I would suggest that such estimates of the amount of missing information are closely related, and probably proportional, to the person's estimate of the difficulty of obtaining that information - ie the larger the perceived difficulty, the larger the size of the perceived lack.

Asking someone to describe the information that they perceive as lacking is a frustrating and exasperating experience for both the questioner and the answerer in all but the most trivial of situations. Given that the information lack is an elusive and unquantifiable concept, it follows that the corresponding information need to fill the gap is equally elusive. It would seem then, that the idea (popular in IR) of a clearly defined, or well understood, information need is at best misleading and at worst ridiculous.

This calls into question the idea that someone can have an identifiable goal or target (in terms of information) to which their information seeking activities are aimed. Only full retrospective knowledge "after the fact" can identify the information that was to ultimately satisfy the information lack. Only then can that information be regarded as the agent's original information need, and thus the target of their information seeking activities.

The point I am making is that it is an information lack (and the perception of such) that is the originator and motivator of the searching activity. An information need is that which is lacking, and that can be said loosely as being sought. This distinction is a subtle one. I introduce it purely to highlight the problems associated with describing something we don't have - where 'not having it' means (in terms of information) not knowing what it is. With the caveat of that subtlety, the two terms can be regarded as referring to the same thing. Given that talking of information needs is traditional in the field of IR, and that it is an arguably less clumsy term, I shall adopt that custom.

Let us imagine a cognitive agent capable of perceiving and acting upon information needs. Let us imagine further that immediately after an information need has been perceived, all flow of information to the perceiving agent stops and all flux of information within that agent ceases. It is obvious that, under such circumstances, the agent will never obtain the missing information. If the restriction on internal flux is lifted, then it is possible that by generating new information internally the agent may derive what was perceived as missing (this ignores the issue of whether, in the absence of external confirmation, the agent would be able to recognise the information as that that was originally missing). Although possible, such purely internal situations are not of particular interest to us as we are concerned primarily with the flow of information into and out of the agent. We wish to optimise that flow such that it leads to effective (in terms of speed and accuracy) resolution of information needs.

With no flow of information into or out of an agent, it will not be able to obtain the information that it requires to resolve its information need. As soon as information begins to flow into the information seeking agent, it will be processed. This processing, whatever its form, will naturally be targeted towards the resolution of the information need. Two clear applications can be seen for the incoming information - firstly, it will be evaluated to see if it is, or is part of, the information being sought, and secondly, it will be incorporated into the general store of information that makes up the agent's view of the world.

The general store of information is a result of all the information that has previously flowed into the agent - ie the agent's previous experiences. Given a different store, the agent would perceive and act differently as a result of a given input of information. With every experience, the store of information will change - even 'identical' experiences (if such a thing were possible) would provoke minimally a reinforcement of a previous pattern. Therefore, the instantaneous state of a store can be regarded as a specification or definition of the agent at that particular time. I will refer to this as simply the current 'state' of the agent.

An agent's perception of an information need will be derived from its state. As this state will be changing as it is exposed to information during its information seeking activities, its perception of the information need will change also.

During seeking activity, an agent will be exposed to much information - a high proportion of which will be similar to, or related to, the information being sought. This proportion will be significantly higher than would be the case during some other arbitrary activity.

An agent will actively be trying to improve its chances of locating the desired information, and so will be most receptive to any such similar or related information encountered along the way. Being 'receptive' means making great effort to incorporate encountered information into its store and using it to both reappraise/redefine its information need, and to process any subsequently encountered information. As a result, we can expect the agent's

perception of its information need to be effected to a greater extent by the changes of its state provoked by these exposures, than by those provoked by other experiences.

From these discussions of 'perceived' information needs, it might be thought that a disembodied, abstract information need exists, and that an agent merely constructs a subjective, dynamic view of it. This would be to misunderstand the nature of the thing. The information need is only present relative to, and contained within, the state of the agent. If the perception of an information need were to somehow spontaneously disappear from the state of an agent, then the information need itself would disappear also. The information need and its perception are one and the same.

Given that the perception of an information need *is* the need, it can be seen that its current and future form is dictated by the perceiving agent's state. Therefore, we could consider the information need as part of that state. That state will change as a result of internal activity and through the agent's perceptions of the outside world. Therefore, to capture the information need fully we would have to include the whole of the agent's state. I would therefore suggest that the agent's state and the information need are one and the same. Therefore, if we talk of changes to the agent's state we imply changes to the information need - and vice versa.

The agent's behaviour (ie responses to external situations) is determined by its state. It seems reasonable to assume that those parts of the state most related to the current activity are those that would most influence the behaviour. It would then be consistent that the information need would affect behaviour related to itself (eg information seeking behaviour) more than other behaviours, and further that it would have the primary effect for that particular behaviour.

As any cognitive activity on the part of the agent will result in changes in its state, any attempt by the agent to describe its currently perceived information need will result in a change of the state of the agent. This change is most likely to be in the areas relating to the information need as those would be the 'thoughts' that would be primarily processed and analysed in the description attempt. The identity between the information need and the agent's state means that the information need itself will be modified by this attempt. This implies that a description attempt will result in a description (regardless of its completeness) that is both inaccurate and immediately outdated.

To summarise:

- An information need is, by its very nature, impossible to describe or define - we cannot describe information we don't have.
- An information need is present only in the mind of the information seeking agent.
- An information need will change with time - as a result of the agent's exposure to information.
- Information encountered during information seeking activities is more likely to change the information need to a larger extent than that encountered elsewhere.
- The behaviour of the agent during information seeking activities is most likely to be influenced to the greatest extent by the information need than by other perceptions.

1.2 Willingness to describe information needs

If we accept the idea presented above that an information need is a constantly changing, inaccessible phenomenon present only in the mind of the searching agent, it becomes clear that it is something that we cannot, and could never, capture. Why is it then, that we are so willing to attempt to describe such things? Further, why do we build computer-based searching systems that rely upon such clearly inaccurate descriptions? The answers to these two questions are directly linked and mutually supporting. The reliance of the vast majority of searching systems upon descriptive queries fuels the expectation and reinforces the belief that this traditional approach is a 'natural' approach. It is natural only in the sense of the approach itself being simple and immediately understandable. This simplicity and understandability is a procedural one - the process of describing (effectively asking for) something that we want is one very familiar to us in our daily lives. I believe that the naturalness of this act is misleading.

In general, we ask for things by communicating a partial description of the desired item to others. This might originally have been done by us as cave-men by pointing and grunting if the item in question was in sight. Requesting items that were not directly in sight or that were more abstract in nature would have been extremely difficult, if not impossible.

Languages developed from the simple grunting and pointing to support these more sophisticated communication needs. Nowadays we find it trivial to request things that are not directly accessible, or things that are purely abstract, by communicating a descriptive request to another person.

I believe that communication is a tool that is applied only when deemed necessary or appropriate - ie when it is not possible to obtain what is desired directly. For example, perceiving a need for a pencil and quickly identifying one upon a nearby desktop will result in the need being satisfied without communication. Communication would be resorted to only when a pencil could not be identified immediately (eg as a result of a clutter of other items), and the cost (in terms of effort, time, and likeliness of failure) of performing a search of the desktop were perceived to be greater than that of formulating an appropriate request and communicating it to an appropriate person (eg turning round and asking an office-mate for one).

What I am saying is that, although we may not realise it, we automatically accept and deal with the fact that there are costs involved in descriptive communication. These costs determine our confidence in obtaining the information we wish. The costs involved when dealing with other humans are relatively low as a result of the degree of common experience and common intention. Descriptive communication between such similar individuals is part of our everyday existence. I believe that, as a result of its everyday nature, we fail to identify the very high expectations that we have of the other person's ability to understand and act upon our communicative requests. Given the enormous difference between humans and computers, such expectations are wholly unrealistic and unreasonable within the context of human- computer communications.

1.3 Salient features of information needs

I will identify the particular characteristics that I regard as key to both understanding the nature of, and providing support for, information needs. These characteristics are considered within the environment of performing a searching task, and as such offer a framework for the evaluation of the support provided by searching systems for an agent's information need. These are not different 'types' of information needs - they are general characteristics of all information needs within a searching environment. I identify three structural characteristics (development, multiplicity, and tangentiality), and two operational characteristics (embedding and threading).

1.3.1 Structural characteristics:

Information needs are developing

The first characteristic I wish to identify is that of the 'developing' or 'evolving' nature of an information need during a search. This was motivated previously, where I described that an information need changes constantly as a result of the exposure of its perceiving agent to information of many sorts - particularly to information closely related to that being sought.

As an agent progresses with its search, its internal representation of the information need (which, as outlined already, is equivalent to the need itself) is augmented by the encountered information. It could be 'augmented' in the sense of the provision of evidence to support or deny beliefs in various aspects of the need - such as the need's importance to the agent and its higher goals, likelihood of satisfaction in general, or in the likelihood of satisfaction in the current environment, etc. It could be augmented simply by the reduction in its 'size' through partial and piecemeal satisfaction. I would argue that this process continues until the information need reaches its most developed state precisely at the instant when it disappears as a result of either satisfaction, or of it no longer having a perceived importance to the agent.

The point is that the information need is not merely changing in a random or arbitrary manner, it is developing in an improving and refining sense - heading inexorably towards an end-point of either satisfaction or redundancy. It is perhaps worth pointing out that this end-point is simply the final state that the information need happens to reach, and it is not in any sense an expected, predictable, or predefined 'final goal' for its existence. Actual end-points can only be discussed in the light of full retrospective knowledge.

In the following discussions I shall use diagrams to assist in the description of the individual characteristics. In them, an information need is presented as a directed line segment - running from the information need's inception, through its development, to its end-point of satisfaction or redundancy [fig 1].

"Inception" is the point at which the information need first affects the state and behaviour of the searching agent.

"Development" is the progression of the information need as it chooses to expose itself to, and is affected by, information. These line segments are presented as straight lines as they represent the history of the information

need's development from its own perspective. This follows from the arguments earlier that information needs exist only in the mind of the perceiving agent. Any other presentation would be to impose some arbitrary external view of the "themes" or "directions" that the information need embodied. This is in stark contrast to, for example, the "Berry picking" view of information needs [2] where the information need is viewed from such an external perspective.

"Satisfaction or redundancy" is the final state of the information need before it ceased to exist as an information need. The arrowhead merely indicates the direction of the development ultimately leading to the point of termination.

It must be kept in mind that these diagrams are presented in the light of full retrospective knowledge - ie after all the information needs have passed from inception, through development to ultimate satisfaction or redundancy.

Fig 1. A developing information need.

Information needs are multiple

Within the context of a particular searching session, it is unlikely that an agent will have a single information need that will be followed exclusively from beginning to end. It is likely that in the course of following a need, other needs will be provoked as a result of exposure to information encountered along the way. In some circumstances these needs may actually be perceived as sub-needs - ie information needs whose satisfaction will contribute to the ultimate satisfaction of one of the other information needs [fig 2].

*Fig 2. Multiple information needs.
(A-D indicate end-points)*

Information needs are tangential

The themes of the additional needs provoked during a searching session are unlikely to be random. Although it is possible that some of the spontaneously provoked needs will not be directly related to the original one, they are most likely to be closely related to that of the current need at the time of their inception. In fact they are most likely to be extremely subtle variations of the information need current at the time. I describe this splitting apart as 'tangential' to highlight two things: firstly, that the two information needs share a common history; and secondly, that they pear apart gently at first with accelerating divergence [fig 3]. This gentle peering apart applies only at the time of the split - very soon afterwards the two needs could be quite distinct. It is intended to highlight that the user may not necessarily perceive the splitting at the point at which it occurs - this may happen only later when the difference has become more apparent.

Fig 3. An information need splitting tangentially into two information needs.

Incorporating tangential splitting into the visual representation gives a tree-like structure of multiple diverging information needs all effectively rooted at the same starting point [fig 4]. In this and later diagrams, line segments will be shown parting from each other. The distance between any points on separate branches has no meaning in this presentation - the separation merely indicates that a distinction exists between them.

*Fig 4. Multiple information needs sharing a common root splitting tangentially from one another.
(A-D indicate end-points)*

1.3.2 Operational characteristics:

Information needs are embedded

If we accept that an information need may fragment into multiple sub-needs, or may provoke a secondary related need at any point during its existence, then the same can happen to those new information needs. This indicates that information needs are naturally embedded within one another. The degree to which they are embedded is most probably bounded by the individual agent's ability to deal with the information necessary to manage the stacking of, and returning to, individual needs.

We can show embedding in the visual representation by means of fragmented lines, with each segment numbered to indicate the order in which it was followed by the agent. For example, *fig 5* shows that the agent followed A for a short distance before being (for example) distracted and following B to completion before returning to follow A again. Segments 2 to 4 show the agent stacking A in order to follow C, and then stacking C in order to follow D, before returning to complete C, and then finally returning to complete the original information need A.

At this point it is perhaps worth reiterating the 'common history' aspect of the tangential relationship between the information needs. In *fig 5* the section marked 1 is the complete development of B, but approximately the first half of that is also the early development of A, C, and D).

*Fig 5. Embedded information needs.
(1-4 indicate order of execution)*

Information needs are threaded

Humans are not naturally the most methodical of creatures. Spontaneous behaviour is both natural and automatic and it is, by its very nature, neither methodical nor disciplined. We develop disciplines in most areas of our life in an effort to gain a certain amount of control and encourage consistency and reliability.

Expecting a searching agent to manage meticulously the numerous and various embedded information needs present during a search session is at best unreasonable, and at worst unworkable. What is much more likely is that the agent will follow one need partially, then switch to another, follow that partially, switch back to the original, follow that partially, before switching once again etc.

These switches of context and sub-task can be provoked by many factors. One such reason might be that of information encountered whilst following up one need provoking a change in a 'stacked' previous need of such a magnitude as to warrant an immediate return to it. Additionally, the reappraisal of a sub- or secondary-need previously regarded as having been followed to satisfaction may be provoked by newly encountered information. This reappraisal in the context of new information may result in the older, satisfied, need being perceived differently and hence regarded, once again, as unsatisfied. That would prompt a return to it in order to develop it further.

I use the term 'threading' to indicate that the execution of each of the information needs is threaded through the execution of the others. It can be thought of as a kind of coarse-grained parallel execution of all the needs from their common route to their individual end points.

Fig 6 indicates that the agent followed *A* before moving off to follow *B* for a short distance, then returned to following *A*, moved off to follow *C* partly before jumping back to continue and complete *B*. The diagram does not indicate why that particular sequence occurred, only that it did.

*Fig 6. Threaded information needs
(1-9 indicate order of execution)*

1.4 Summary

In this section I have attempted to show: that describing information needs is an ill-founded and ineffective endeavour; that information needs change constantly and that this change is most affected by related information encountered during information seeking activities; that the behaviour of an agent during searching activities will be greatly influenced by the current state of the information need.

I then described characteristics of the structure of information needs during a searching session: the constant change of an information need is a development towards an end point of either satisfaction or redundancy; the original information need may split into several sub or secondary needs, each of which is a different development of

the original need leading up to the split. I said nothing about how the formation of the structural characteristics happens, simply that it does.

Finally I identified two related characteristics of the execution of the information needs by the searching agent: the development of an individual information need may be embedded within the development of another; the embedding might not be a clean completion of individual information needs, but may result in a threading of the development of the information needs through each other. These are essentially a recognition of the operational characteristics of information needs and that these characteristics are managed by the agent.

2 Current IR systems

In this report I consider only interactive IR systems. Batch systems are so hopelessly far removed in their mode of operation from the realities of the searching task that evaluation, in terms of support for information needs, becomes so subtle as to render it practically meaningless.

It is worth noting that it is primarily the user-interfaces of the interactive systems with which I am concerned - ie how a system presents the information objects, and its retrieval techniques for searching through them, to the user.

Terms of reference

I will list briefly the terms of reference that I shall use to lead the discussions in this and the following sections regarding the provision of support for information needs in IR systems. These result from the discussions in the previous section.

Avoidance of:

- description of information needs

Support for structural characteristics:

- development
- tangentiality
- multiplicity

Support for operational characteristics:

- embedding
- threading

The wider context within which I am carrying out this work is that of multimedia IR where the goal is the development of a new framework for IR that will facilitate effective retrieval of information expressed in the new media of sound, image, video etc. At this point in the discussion there is no clear connection between that goal and the purely information need related terms of reference. Its significance will become apparent in the following discussions. For the moment I simply add it as a term of reference:

Support or suitability for:

- the new media

Query-based and browse-based systems

Within interactive approaches to IR systems I see two main classes. Both classes represent a fundamentally different approach to exploration of an information space by the user - that of query-based and that of browse-based. This division corresponds, in effect, to my first term of reference - that of the reliance upon describing information needs.

Systems employing the query-based approach rely upon some kind of description being provided and refined by the user. Systems employing the browse-based approach rely upon the user selecting or pointing to information objects presented to them by the system. I see this distinction as key to understanding where we have been going wrong in the past and how we can take the interfaces of IR systems onto the next generation. In this section I will

present my views, first on some of the current query-based systems, and then on the current browse-based systems in an effort to encourage that understanding.

2.1 The query-based approach

I regard as a query-based system any one that requires of the user the production and refinement of a description of their information need. Broadly speaking, this description encompasses the vast majority of current IR systems - both commercial and research.

The most prominent examples of query-based systems are the Boolean systems offered by commercial dial-up online information providers - eg *STAIRS* [6], and its derivative *BRS-Search*. Other Boolean systems available for most personal workstations and bundled with information bases on CD-ROM are: *Topic* [10], *Personal Librarian*, and the rather dismal searching interfaces provided with almost every non-IR computer system. The most modern of the query-based systems are the weighted systems that unfortunately are generally restricted to the research world: *ConQuest* [12], *NRT* [18], *WAIS* [20], *Inquiry* [4], *Smart* [17], *Okapi* [16].

What all these systems have in common is that their operation centres around a query. A query is the description of the user's information need, presented in terms of some system supported language. These languages are direct derivatives of the internal representations used by the system for the stored information objects. For example, the presence or absence of individual words is the representation of documents used in the *STAIRS* and *BRS* systems, and that, along with Boolean operators such as 'AND' and 'OR', make up their query languages.

In practice, the user generates an initial query, the query is evaluated by the system, a list of resulting documents is presented to the user, the user inspects these documents, refines his query, the new query is evaluated by the system, and the process continues iteratively [fig 7].

Fig 7. The iterative process of query-based systems.

Reliance upon descriptive searching

Users must describe the documents they believe will satisfy their information needs by describing the pattern of presence or absence of words within them. This must be done without having seen the documents in advance. As I have pointed out already, expecting a user to be able to describe something he doesn't have is not a very promising approach. Placing the restriction that such descriptions must be made in terms of an artificial language, is to further hinder a user satisfying his needs. These languages are chosen mainly for their computational simplicity and correspondence to the internal representation schemes and not for their suitability for expression of information needs.

Weighted retrieval systems (eg *Smart* and *NRT*), have slightly different internal representations and significantly different retrieval algorithms to that of the Boolean systems. This results in querying languages with very similar bases (ie word occurrences), but with very different interpretations. Here, the 'hard' nature of Boolean operators on word occurrences is replaced by 'soft' weights attached to candidate words giving more informal, less error-sensitive languages in which to express information needs. Nevertheless, the direct correspondence between the internal representations and the query languages remain.

These representations are also used by the systems to 'explain' particular retrieval decisions. The words that caused the retrieval of particular documents are often highlighted in some way, thus providing feedback to the user on the extent to which each of the query words is involved in the retrieval of documents. This explanation is crucial to the effective exploitation of such systems. It is with this information that users make their decisions upon how to refine their query and thus continue their search.

Support for developing information needs

During the query-refinement/query-execution/document-inspection iterations, the user's information need will be changing as a result of his exposure to documents. One could imagine that the user would be able to reflect that changing information need through a corresponding modification of the query. Nevertheless, two things prevent this from happening.

The first factor is that as an information need changes, the changes might be sufficiently subtle that the user would not be consciously aware of each change, despite that it had, for example, already affecting his relevance judgments. Changes in our beliefs or attitudes often are not clear to us before we notice that they are having significant effects upon our actions. A user will only become aware of a change in his information need after it has reached a level great enough to provoke a reappraisal of that need. Further, time will have passed between the changes taking effect and the user becoming aware and then acting upon them. The relevance judgments being made and the user's general satisfaction with the progress of the search will naturally be out of synchronisation with the way that the search is being driven.

The second factor is that all information needs perceived by the user must be described. This requires that he translate each new version of the information need into the terms of the query language. This task is not simple and the feedback provided by the results of execution of a particular query are the only way in which a user can appraise the appropriateness of a particular query formulation. This means that several iterations of query-refinement and query-execution are likely to be required before the user will be happy with, or confident in, the query as a representation of his information need. Nevertheless, in the meantime, the very information need that he is endeavouring to represent will have changed 'beneath his feet'.

Query-based systems do not *support* developing information needs. These systems *allow* information needs to develop and the user to follow that development. Nevertheless, the requirement for descriptions serves only to hinder this development. What they do support the development of, are *queries* - but as indicated above, they are a poor surrogate for the information needs themselves.

Support for multiple information needs

The presence of multiple information needs during a search session implies that the user will want to switch between those needs at any point, any number of times. Query-based systems, with their emphasis on the progressive development of a query naturally hinder this process. As outlined above, proper development of a query can only happen with the query-refinement/query-execution/document-appraisal/query-refinement cycle. This means that either the process of effective query development is disturbed, or that the new information need cannot be followed up as and when necessary.

An additional problem is that we might not even notice that we have had a new, or subordinate, information need provoked until we have already followed it for some time. This effect is directly related to that outlined above in the discussion of the support for developing information needs, where changes in an information need is not consciously perceived until it has become relatively large. Query-based systems with their iterative query-refinement do not support directly the convenient 'unwinding' of the previous actions.

Histories (or saved states) of all queries and their resulting documents lists can go some way to supporting this. Nevertheless, histories would typically store only those queries submitted for execution - they would not store each of the various edited versions that were present between query-executions. Even if that level of detail were stored, it would result in inordinately long and complex histories - histories that would only show the results of the attempts to represent instantaneous information needs. That would be rather indirect. Much better would be the information that actually provoked the changes in the needs in the first place, as that would be much more likely to provoke a similar response upon reappraisal.

Support for embedded and threaded information needs

The fact that multiple information needs are embedded and threaded, is basically the result of the manner in which the user spontaneously perceives of, and manages the pursuance of, those needs. With the rather dismal lack of support already outlined for multiple information needs in general, it is pointless to discuss the much greater, and very specific, support that threading requires.

Support for the new information media

Query-based systems have been the major thrust of IR research for many years, and look set to remain so. The concept of a query has become central to our ideas of IR systems. Despite the success of these systems in the past, I believe their time has been and gone. The hardware and software tools to support the generation and storage of multimedia information have become widespread. Large collections of such information are gathering at an increasing pace, and the desire to retrieve that information is growing in parallel.

Information stored in a non-textual medium, by its very nature, requires different representation techniques to that of text. Examples like that of speech-only audio information being transcribed into text, and then that text being used for retrieval of the original audio are, by virtue of their triviality, not very instructive and are simply unrepresentative. Music, photographs, video, etc cannot be given anything but the most trivial of textual 'transcriptions'. I believe that we will, by necessity, develop new representation techniques for non-textual media. These representations, if they are to be effective, are unlikely to be accessible to humans. One such approach might be that of using neural networks, where any representation would consist of minute probability modifications spread across many thousands of cells each of which has no 'real' meaning to us. As a result, I believe the reliance of query-based systems upon the user's exposure to, and manipulation of, internal representations, simply rules out the descriptive approach as the basis of the next generation of multimedia IR systems.

2.2 The browse-based approach

In a browse-based system there is no query, and no explicit description of the user's information need. The user furthers his search purely through selections or indications over information objects (usually documents) presented to him by the system. Such systems almost exclusively employ a graphical interface presenting documents as objects with links between them [fig 8].

The classic browse-based IR system is I^3R [22] - where documents and groups of documents are presented as icons with accessibility links between them. *Transient-Hypergraphs* [19] exemplify the application of the browse-based approach to a traditional strongly typed database. *TACHIR* [1] is a system that allows the user to browse across a plane of documents and their accessibility relationships, across a similar plane of terms, and across the accessibility relationships between the different planes. *Galaxy of News* [15], and *Bead* [5], present documents in a roughly 2-D space that has been dimensionally reduced from an original higher dimensional space. The distance-metric, is again, that of an accessibility relationship.

Fig 8. A network of information objects, with accessibility relationships between them

One might ask what is special about these systems with respect to hypertext systems. I make the distinction based upon the origin and nature of the primary accessibility relationships involved. In hypertext systems the accessibility links are either human generated, or derived directly from the structure/syntax of the objects concerned. The links commonly span a tiny subset of the object space (ie the degree of mutual interconnection is very low), and the links are of a binary nature - ie present or absent. In browse-based IR systems the links are computed, and that computation is based upon the documents' semantic content (ie not just their syntax). The links are commonly present between every object and every other object in the space. Finally, the strength of the links (ie the degree of accessibility) is represented by continuous values.

Selections instead of descriptions

In order to move around within the information space, the user makes selections from a list of alternatives presented to him by the system. These alternatives are information objects (usually documents) that the system has computed to be the most appropriate next steps. These next steps are computed relative to the current location, ie current information object, the degree of appropriateness being measured by the accessibility relationship defined for the space. This accessibility relationship may be anything from a simple mutual set membership, to a sophisticated measure of semantic similarity between two objects. Usually it is a pseudo-semantic similarity measure based upon one of the classical retrieval models. Although it is possible for every other information object to have some degree of accessibility from a given object, it is not practical, in all but the most trivially small collections of objects, to present them all to the user. Only a small number of the strongest links is presented at each object.

Static non-adaptive views

At this point it would be nice to be able to say that the systems interpret every action made by the user and use this information to update the system-view of what the user is looking for - thus rearranging the space to suit the user. Unfortunately, this is not the case. The information spaces (ie the arrangement of information objects and their accessibility relationships) are static. They are set up once for a particular collection of information objects, and remain the same for all searching sessions for all users. The selections made by the user only allow them to move from one statically linked object to the next. I regard the static nature of the information space to be less attractive than that of the adaptive and dynamic space of the query-based systems where the space adapts to individual users and their instantaneous needs.

During a session, the user simply moves from one object to the next as the feeling takes him. He need only identify something of interest in the locality of his current object and investigate. There is no need for the user to attempt to translate that interest into a description of why he now wishes to move in that direction. More importantly, there is no need for the user to even think in which direction he may possibly wish to move - he simply sees where he wishes to go and goes there. The manner by which information needs are followed is both spontaneous and immediate.

Support for developing information needs

As the information space is static, the paths through it can only be over the predetermined links between information objects. The small number of statically defined links available at each node restrict the possible next steps to those that seem appropriate from the point of view of the current node alone. The 'view' from that object is the same for all information needs, regardless of their state or rate of development.

To assist in appreciating this point, let us imagine two users, identical in every respect and therefore with an identical initial information need. If they were to approach a particular object by different routes, they would have been exposed to different information (or minimally the same information, but in a different order). This would imply that they, and their information needs, at the time they reached the object in question, would be different. In that case their 'view' of the world would be different and the effect, on their information needs, of their exposure to the information object would be different. This implies that to support this, a system should somehow provide different candidate objects and links for these two users - static browsing systems do not, and thus do not support developing information needs.

Support for multiple information needs

The almost exclusively graphical user-interfaces of these systems provide the well accepted benefits in visibility, immediacy, and ease of management of the objects being presented. In browsing systems, it is the information objects and their links that are presented.

The graphical presentation of objects and links means that no matter what route is taken through the object space, that route is visible and can be backtracked or retraced at will. Additionally, it is the actual objects that are presented in these 'histories', not user descriptions of information needs as is the case in query-based systems. It is those objects that provoke the changes in the information need in the first instance and so they should have a direct effect of reorienting the user if they were to be returned to at a later stage.

Depending upon the particular system, various amounts of the information space will be visible at any one time. The more of the space that is visible, the more, in general, the user will be able to simply jump to a previous location without actually having to backtrack through every intermediate step.

This visibility of information objects and the orientation provided by the links between them provide an environment that supports the management (by the user) of multiple information needs.

Support for embedded and threaded information needs

As outlined above, switching between paths through the information space is supported by the graphical presentation. This allows users to follow an information need, and then return to a previous one by simply backtracking until they identify the information object at which they feel they began to diverge from the primary and follow a secondary information need. They can then return to following the original primary need. Embedded needs are therefore supported.

Taking this one step further, it is easy to see that it is not necessary for the user to follow one particular information need to completion before backtracking to continue a previous need. They can switch at will between several concurrent information needs, thus threading the development of them.

Although there is support for multiple, embedded, and threaded information needs, it must be remembered that the user's information needs, however they are developed, arranged, or managed, must fit into the restricted scope provided by the static information spaces of these systems.

Support for the new information media

In browse-based systems the user is not exposed to the internal representations used to store information objects or compute accessibility relationships between them. The user simply views the objects in their native format, be it text, sound, graphics, etc. This means that the most inaccessible, complex, and contrived of combinations of mechanisms can be employed to structure the information space. I regard this as an attractive feature of browse-based systems, particularly when considering the development of a new generation of multimedia IR systems.

2.3 Considering a hybrid approach

From the critique of the two main approaches to IR it can be seen that the support, with respect to my terms of reference, is rather patchy. Nevertheless, I see a pattern in this support. The query-based systems' strong point is the highly adaptive nature of the views of the information space enabling them to follow the changing information needs. Their main flaw is the difficulty in driving this adaptability effectively resulting from their reliance upon descriptive queries. The browse-based systems' advantage, resulting from their graphical interface, is the highly direct and immediate manner in which a user can act upon observed information and changes in his information need. Their drawback is the non-adaptive, unsympathetic, static views of the information space.

Graphical presentation of the information space

The graphical environment of information objects and links between them is the tool provided by browse-based systems that supports the multiple, embedded, and threaded nature of information needs. The network of objects and links provides an overview of the current search path. This overview provides orientation-, coordination-, and

memory-aides that allow the user to follow his unpredictable and undisciplined information needs. For example, if the user decides that his current direction has become unsatisfactory, he can either iteratively backtrack until he is on firmer ground and head off in another direction, alternatively, he can immediately jump back to a point far back where he is certain he was doing better, and continue more carefully.

Adaptive views of the information space

Query-based systems restructure the information space on each query evaluation. This is a method by which the view presented to a user situated at any point in the space can be made to correspond more closely to, or at least change in a manner sympathetic to, the user's current information need. The system attempts to restructure the space in a manner that brings closer, and thus makes more accessible, those information objects it considers to be most likely to be relevant to the user's instantaneous information need. In principle, adaptive views, support the developing nature of information needs. Nevertheless, the success of such a technique would be contingent upon the system obtaining an accurate representation of the information need. As stated previously, the user-generated descriptions of query-based systems are far from ideal representations.

In the browse-based approach, there is a clear sense of a user always being situated at one or other information object in the information space. In the query-based approach the user is never at an actual object in the space. They are at a virtual point, the position of which is defined by the query. That point may happen to correspond to the position of an information object (eg when the query contains exactly the same information as an object in the space), but the user is never considered to actually *be* at the object. To use the suggested graphical presentation technique of objects and links we must find a way to provide adaptive restructuring of the space, but from the perspective of location at individual information objects.

'Black box' representations

I believe the key to supporting the retrieval of complex media objects is to remove all traces of the system's internal representation methods from all interactions with the user. The upshot of this is to effectively rule out all forms of descriptive querying. I find it impossible to imagine a language that would allow a direct description of something to be formulated without resort to some representation technique. Even when we attempt to describe abstract concepts or visions to other people we must employ a basic representational alphabet. Such an alphabet must often be negotiated in advance of the description being given because we all have different, but commonly structured, views of the world and its contents. This once again highlights the mismatch between the knowledge and experience of a human and that of a computer system.

Summary

View adaptation by space restructuring appears to be a promising technique for the support of the developing nature of information needs. To drive that adaptation/restructuring we must obtain a representation of the user's current information need. Unfortunately, in order to support the new media (amongst other reasons) we cannot use direct descriptions of the information need from the user. Therefore, a new way of obtaining a usable representation of an information need is required.

In order to support multiple, embedded, and threaded information needs we have the graphical presentation of an information space. Unfortunately, this appears to be at odds with the desire to provide an adaptive restructuring space. Therefore, a method of providing a graphical presentation of a changing information space (without resulting in user disorientation) is required.

3 The new, ostensive approach

In this section I develop and describe the fundamentals of my proposed new approach to interaction between IR systems and their information seeking users. I refer to it as an 'ostensive' approach because it exploits ostensive definitions made by the user of his information need by mechanisms such as implicit relevance indications.

The Oxford English Dictionary [13] defines an ostensive definition as follows (senses 1a to 2a omitted):

2. b. *ostensive definition* (Philos.), the explanation of a word by presenting, pointing at, or otherwise indicating one or more objects to which it applies.

A classic example of an ostensive definition is a parent teaching a colour to a child. The parent will point to or present several objects of the colour in question. Over time (ie after several such indications) the child will form an idea of what it means for something to be of that colour. This definition is a perfectly usable one (and one carried on to later life), but it is one that was obtained without any ‘description’ taking place. Further, an understanding of the physical nature of colour is neither being used nor transferred.

The new approach centres around the provision of support for developing information needs.

3.1 Supporting developing information needs

A developing information need is a changing need. The need changes in an improving manner towards an end point of satisfaction or redundancy. This improvement results from the searching agent being exposed to information during the search. This information adds to and refines the knowledge and view of the world that is held by the agent. It is that knowledge that drives the assessment of information with respect to the information need. Therefore, as the information upon which such relevance assessments are made improves, it seems reasonable to assume that the assessments themselves improve. That is, relevance assessments and relevance-related decisions improve during the progress of a search.

The passage of experience

I refer to the above change in the knowledge and view of the world of the agent, resulting from the exposure to information, as the ‘passage of experience’.

If relevance assessments improve as a result of the passage of experience, then this implies that at any point in a search, any previously made assessments are less good - or more specifically, have retrospectively *become* less good. The corollary of assessments improving continually as the search progresses, is that they would be seen to continually degrade if one were to look further and further backwards from any point.

It is now clear why browse-based systems do not support developing information needs. Decisions made at the start of any search path restrict the portion of the information space that is directly accessible from then on. Therefore, at any later point in the search, the available space is still being restricted by those early and now less-appropriate assessments. One solution might be to completely reappraise the early decisions in the light of the passage of experience - ie start all over again.

Undoing the effects of early decisions

Instead of the rather drastic action of starting again, it would seem useful if the effects of the earlier (and thus less appropriate) decisions could be removed from the current ‘view’ of the information space. One could consider an approach such as a ‘sell-by-date’ that would simply remove completely all decisions (or more specifically, the effects of decisions) made over a certain amount of time in the past.

The idea of ‘time’ would not necessarily be measured in minutes and seconds, it would most likely be an ‘information time’, where time progresses one unit at each atomic flux event of the information need. Such events could be, for example, the exposure of the searching agent to a new information object. It is these exposures that provoke the changes in the information need, and if there has been no change in the information need then there is no point in changing the degree of appropriateness of a decision or assessment. The more exposures, the more changes. Therefore, a reasonable ‘tick’ in this time would be an individual presentation of an information object to the searching agent.

Expiring old decisions and assessments completely after some time interval is itself rather drastic. It also does not accurately reflect their range of appropriateness to the information need. All decisions and assessments have some degree of relevance to the information need. Rather than simply removing completely the effects of old decisions, a gradual reduction of their effect over time would correspond more closely to their validity [fig 9].

Fig 9. Two examples of functions to reduce the effect or 'weight' of ostensive evidence with its age.

If the effect of decisions is to be reduced gradually over time, a manner of determining the 'age' of particular decisions with respect to each other and with respect to the particular information needs is required. I propose that the path of information object leading to the current point be used to order, and thus age, the decisions. There would be no absolute sense of age - only a relative one. This path would be from the point representing the start of the information need, and would pass through all information objects that contributed to the state of the current information need. One could consider this path as a history and ostensive definition of the information need.

To capture the ostensive definition of the information need, the system would gather all evidence available along the path. Such evidence might include the information objects to which the searching agent was exposed, information objects regarded as relevant by the agent, and decisions or selections made by the agent. The evidence most ostensive in nature would be information indicated by the searching agent, whether directly or indirectly, to be relevant. These need not be information objects themselves, they could be decisions or actions motivated by the agent's perception of relevance in a more general sense.

The ostensive definition collected by the system would be used to restructure the information space around the point at which the searching agent is currently located. That is, the view presented to the agent will be calculated using the ostensive definition. This is essentially the same as in the query-based approach, except for the source of the representation of the information need - ie an ostensive definition rather than a description generated by the agent.

3.2 Contextual interpretation

The view presented at any point in the information space will consist of other information objects. These would ideally be those objects that the searching agent would most likely want to investigate after exposure to the current object. The browse-based approach presents objects that are most similar to the current one (the particular definition of 'similar' is detail that does not concern us at this level). This calculation of similarity is a pairwise comparison of the objects, one of which is the current object. I consider the objects as being *interpreted* by the similarity measure. Nevertheless, this interpretation is being performed purely within the context of the two objects concerned, or at best, including some summary information on the profile of the wider information space. This localised and static context within which the interpretations are made does not reflect the nature of the context within which the searching agent will make interpretations.

I find it impossible to imagine any information, however 'simple', that could not be interpreted differently in different contexts. Information objects will have many (potentially infinite) interpretations and so restricting the view presented at these objects to one static selection of next steps seems unreasonable.

In the approach that I am proposing, the view presented at any point will not simply depend upon the point itself, but will depend also upon the path taken by the agent to reach that point. The idea is that if the searching agent is following up a particular theme, that theme will be ostensively apparent in the path in the form of such things as the selected information objects and their order of being visited. If the path were used to determine the view, the view could then reflect that theme. The path would be effectively the context within which the information object would be interpreted by the system.

The effectiveness of this approach is crucially dependent upon the manner in which the evidence collected along the path is used to restructure the information space. It is impossible for the system to have a full and accurate copy of the context used by the searching agent to make its interpretations. Nevertheless, it is hoped that with appropriate methods of evidence identification, collection, and weighting a system context can be obtained from a path and made to at least approximate the general themes and directions taken by the searching agent.

The view from an information object will depend upon the path taken to reach it [fig 10]. This implies that if a searching agent were to revisit a particular object (either because two different paths lead to the same object, or because an object may be present at more than one point on the same path) they would see a different view. This might suggest that any particular object may be present in multiple locations in the space. This seems inconsistent with the common understanding of a space - ie no matter how one gets to a point, it will be the same point, in the same place, and therefore have the same view around it.

Fig 10. Two different paths to the same object and their different views of the candidate next steps.

In the proposed approach, each object can be regarded as a source of multiple interpretations - each of which is a result of the same object and different contexts within which the interpretations are made. The contexts are the paths taken to reach an object. It is the interpretations, not the actual information objects, that make up the points in the space through which the agent moves. The agent is therefore moving in an interpretation-space. The interpretations are unique and thus the apparent inconsistencies are no longer present.

Problems may arise when trying to present this space graphically because it is possible that real users (ie not the idealised “agents” of which I have spoken so far) may have difficulty in separating the idea of a document from the idea of its interpretation. As a result, they may be confused or disoriented by the apparent reappearance of the same document in several places in the space, each time having a different local space around it. Experimental evaluation may well be the only way to quantify this effect. Techniques may have to be developed to support, encourage, and reinforce the idea of browsing interpretations and the secondary nature of simple information object identity.

3.3 Summary

- Visual browsing environment of objects (actually, interpretations) and links.
- Objects in the space are actually interpretations of the underlying information objects. Therefore, individual information objects may appear in multiple locations in the guise of different interpretations.
- Small number of objects presented at each object as the most likely next steps.
- Selection of the most likely next steps based upon a contextual interpretation of the current base-object.
- Contextual interpretations made using an ostensive definition of the developing information need.
- Ostensive definitions obtained from evidence collected along the path to the current object.
- Evidence weighted by degree of uncertainty related to its ‘age’ in the path.

- Evidence made up of things like implicit relevance assessments, decisions made, and information to which the agent was exposed.

The ostensive approach can be summarised generally as one where the system builds a representation of the user's information need by observing and interpreting the user's actions in trying to satisfy that need. The system then restructures and presents the information space in a manner that is consistent with its current idea of the user's developing information need.

4 Proposed implementation

In this section I give an overview of how the currently envisaged IR system would present to the user the ostensive support for developing information needs. By making concrete the visual interaction environment I hope to provide a clearer idea of how the ostensive approach will work in practice. In addition, I hope that it will facilitate a better overall appreciation of what the ostensive definition approach has to offer.

4.1 Objects and links

The visual presentation will be that of a surface with objects and links placed upon it. Objects will be presented as icons, the exact form of which may indicate some aspect of the information or its medium - for the purposes of this discussion they will be simple document icons. Links will be presented as arrows indicating the direction of movement. Movement will generally be from left to right. A path will be a sequence of objects and links.

There will be the concept of the current object. A small number of candidate next-steps will be shown around the current object. Candidates will only be shown for the current object - thus reducing considerably the visual clutter, with only those objects actually visited being shown on a path [fig 11].

The candidates presented will be chosen by the system based upon evidence collected from the path - ie they represent the view at that point. This allows the view to follow the developing nature of the information need.

Fig 11. A three document path, with five candidate next-steps surrounding the 'current object'.

Backtracking to a previously visited object (ie making it the current object) will show its next-steps again. This would occur, for example, if at some point more than one interesting object was seen and the user wanted to investigate them both - he might follow one for some distance and then return to follow others [fig 12]. This is how new, secondary information needs would be followed up, or alternatively how a user would backtrack to return to the primary information need from which he had been diverted.

Fig 12. Backtracking to review the other candidates at a previously visited object

Each branch in the tree of paths represents a separate information need. Each information need can be thus be seen to have a common heritage. This represents the tangentially diverting sub- and secondary- information needs [fig 13].

Fig 13. Four information needs represented by four path branches.

The user can at any time jump from object to object, review the candidates at that point and then investigate them further. This provides the support for multiple, embedded and threaded information needs.

4.2 Avoiding clutter within a consistent layout

During a complex or thorough search the user will generate many branches. Techniques will have to be developed to ensure that there is always sufficient display space for the new branching, without compromising the utility of the existing layout for orientation. The maintenance of the general layout is key to providing orientation for the user, or more specifically, allowing the user to orient himself - and thus essential for supporting the embedding and threading of information needs.

One technique I propose is that of effectively 'tearing' the surface of the plane upon which the objects are drawn and 'folding' sets of branches underneath others. This would allow another set of branches to grow unrestrictedly on top. Instead of tearing the surface, a depth cue could be used to give the impression that the folded branches are further back into the display surface [fig 14].

Fig 14. Folding (by depth cue) of one set of branches under another to allow growth.

By applying a spherical transformation (known as a “fisheye” view [9]) to the display surface, the parts of the path-tree that are distant from the current object are reduced in size [fig 15]. This means that the clutter is effectively brushed aside, providing a clear unhindered view of the area of immediate interest. The overall visual context and its orientation effect is still present nonetheless.

Fig 15. A fisheye view of the display surface.

Using this technique, it might be possible to avoid folding and tearing. As the current search path expands, the surrounding area could simply be pushed out of the way. This would involve modifying the display transformation function dynamically. Additionally, as the branches expand as they go to the right, the complete tree will take on a triangular form - narrow at the left and wide at the right. A simple spherical transformation would clearly not be sufficient if the display space were to be used most efficiently to present such a structure.

A dynamic fish-eye transformation would also allow the individual branches/paths to be further spread apart to allow growth and the resulting spatially impossible spreading to be hidden within the overall spherical distortion. Only the path leading to the current object (ie the part to the left of the current object) would need to be shown accurately, or at least in a manner representative of its shape. All this requires that a highly specialised and adaptive transformation be developed.

4.3 Displaying object contents

At this point I shall introduce the general look of the proposed interface. The display would consist of two windows, a “browse” window and an “object” window - the first would be a view of the display surface, and the second a view of the contents of an object [fig 16].

Fig 16. Separate “browse” and “object” windows.

The object contents window would continually and automatically show the full contents of the current object (or more precisely, the contents of the basic object, one of whose interpretations is the current object on the display). The current object would always be in the centre of the browse window, and be the one with its candidate next steps displayed around it.

There is clearly insufficient space available in the browse window to allow the titles or abstracts of all objects to be shown, never mind their full contents. I propose two things: firstly, moving the mouse over an object (ie “pointing”) to view its abstracted form; and secondly, clicking on an object to make it the current object and thus display its full contents.

Pointing to view an object’s header

Simply moving the mouse over any object in the browse would result in a small transient pop-up window appearing over that object. This window would remain displayed only as long as the mouse was over the object, and would contain a highly abstracted form of the object. This abstract would depend upon the object in question and might be, for example, the title of the document, a miniature version of the image, a short sound-bite from an audio recording, a fragment of video etc [figs 17, 18].

Fig 17. A text abstract being displayed on top of a candidate object over which the mouse is situated.

Fig 18. An image abstract being displayed on top of a candidate object over which the mouse is situated.

Pointing to view abstracts is not restricted to candidate next steps, it applies to any object visible in the browse window. This would allow the user to quickly review the objects in the path leading to the current object, or to remind himself of the objects encountered in other paths [fig 19].

Fig 19. An abstract being displayed over an object in another path.

Clicking to access an object and its contents

When the user has identified an object from its abstract that he wishes to inspect more closely (ie view its full object contents), he would simply click on it. The spherical surface presented in the browse window would then rotate to place the new current object under the centre of the browse window, the object's candidate next-steps would be shown around it, and its contents would be displayed in the object window.

These operations and the behaviour of the system would be independent of the media type of the objects. If the current object is a text object then its text would be shown in the object window, if it were a graphical object, its graphical contents would be shown in the window [fig 20], similarly for other single or mixed media types.

Fig 20. A graphical object as the current object.

The combination of pointing to view and clicking to select should allow the user to flick around the browsing paths at will, quickly and conveniently calling up the desired degree of information relating to an individual object, and subsequently act upon that information without the requirement to use or remember complex commands.

4.4 Summary

I hope that this section has given a 'feel' for how the system might work. The points that I regard as important at this level are that the user need simply 'follow his nose' when looking for something. There is no need for him to stop and think about what he is looking for and how he might describe it. There is no exposure of the user to the representations used by the underlying retrieval systems to generate the views. The user is only exposed to the information objects themselves in the original media in which they were intended for viewing.

I believe that this media independence brings with it a language and domain independence, that is unique to this system. There is basically no dialogue between the system and the user. The user performs his task by performing actions and the system effectively looks on and attempts to assist in a sympathetic manner. I feel this approach could provide an environment in which the user could 'sniff' around, perhaps unaware of a system performing any kind of searching in the background - ie the user might have the impression that he was in fact browsing a static surface of information objects arranged previously into useful thematic trails by some helpful and smart librarian.

At a superficial level, the proposed interface lacks conventional dialogue-boxes, buttons, scrollbars and use of the keyboard. I would hope that this may reduce to a minimum the amount of "computer-literacy" required of a user, and would reduce for all users the amount of procedural operation required to perform searching tasks.

5 Continuing the work

This work is at a relatively early stage and this report is of the initial informal ideas. Nevertheless, a clear direction has been determined, and an abstract framework established. It is the formalisms and implementations of that framework in terms of existing theories, such as the probabilistic retrieval model, that form the thrust of the current work.

Amongst the current tasks are:

- Formalising the hypotheses, arguments, and assertions.

eg. Standardising the arguments in terms of ‘degrees of belief’ and ‘degrees of uncertainty’ instead of the mix of rather vague usages of terms such as appropriateness and validity.

- Including into the discussions direct connections to related work. For example, McKay’s work on ‘states of conditional readiness’ [11], Lucy Suchman’s ‘partial route plans’ [21], and Marcia Bates’ ‘berry picking’ [2] ideas.
- Contrasting the user model implicit in this approach with the significantly more complex and sophisticated models of cognition for IR developed by Ingwersen [7] and Belkin [3].
- Developing a probabilistic method for collecting and assigning degrees of uncertainty to evidence, and then using that information to determine views. Using Ramsay’s ideas on probabilities and preferences [in 8 and originally in 14], and Jeffrey’s work of probability kinematics [8].
- Building a prototype interface to allow the development of (and evaluation of) various display transformations.
- Considering for which multimedia information seeking tasks the approach/system may be best and least suited in preparation for comprehensive user testing.

6 References

1. Maristella Agosti, Massimo Melucci, and Fabio Crestani. TACHIR: a tool for the automatic construction of hypertexts for Information Retrieval. In: proc. of RIAO 94, Rockefeller University, USA, 1994, pp. 338-357
2. Marcia J. Bates. The design of browsing and berry picking techniques for the online search interface. Online Review, 1989, vol. 13, no. 5, pp. 407-424
3. Nick Belkin, R. Oddy, and H. Brooks. ASK for Information Retrieval, parts I & II. J. Doc, 1982, vol. 38, pp. 61-71 & pp. 145-164
4. J. P. Callam, W. B. Croft, and S. M. Harding. The INQUERY retrieval system. In: proc. of 3rd Int. Conference on Database and Expert Systems Applications. 1992, pp. 78-83
5. Matthew Chalmers, and Paul Chitson. Bead: explorations in information visualisation. In: proc. of ACM-SIGIR, 1992, pp. 330-337.
6. Storage and Information Retrieval system (STAIRS) - general information manual. 2nd edition, IBM Germany, Stuttgart, 1972
7. Peter Ingwersen. Information Retrieval Interaction. Taylor Graham, London, 1992, pp. 203-222
8. Richard C. Jeffrey. The logic of decision. McGraw-Hill, 1965
9. John Lamping, and Ramano Rao. Laying out and visualising large trees using hyperbolic space. In: proc. of ACM-UIST '94, 1994, pp. 13-14.
10. John W. Lehman, and Clifford A. Reid. Knowledge-based searching with TOPIC. In: proc. of The second Text Retrieval Conference (TREC-2), 1994, pp. 209-221
11. Donald M. MacKay. Information mechanism and meaning. The M.I.T. Press, 1969
12. Paul E. Nelson. The ConQuest system. In: proc. of The second Text Retrieval Conference (TREC-2), 1994, pp. 265-270
13. Oxford English Dictionary. CD-ROM, 2nd Edition, Oxford University Press, 1993

14. Frank P. Ramsay. Truth and probability. In: "The foundations of Mathematics and other logical essays". Kegan Paul, 1931
15. Earl Rennison. Galaxy of news. An approach to visualising and understanding expansive news landscapes. In: proc. of ACM-UIST '94, 1994,
16. Robertson, et al. Okapi at TREC. In: proc. of The first Text Retrieval Conference (TREC-1), 1993, pp. 21-30
17. Salton (editor). The SMART retrieval system - experiments in automatic document processing. Prentice-Hall, 1971
18. M. Sanderson & C.J. van Rijsbergen, 1991. NRT: news retrieval tool. Electronic Publishing origination dissemination and design, vol. 4, no. 4, pp. 205-217
19. Michael A. Shepherd and Carolyn Watters. Virtual structures for hypertext. In: proc. of ACM SIGIR / ISS International Conference of Multimedia Information Systems, pp. 201-219, 1991
20. R. M. Stein. Browsing through the terabytes. Byte, pp. 157-164
21. Lucy Suchman. Plans and situated action. Cambridge University Press, 1987
22. Rodger Howard Thompson. The design and implementation of an intelligent interface for Information Retrieval. Ph.D., University of Massachusetts, 1989.
23. Frederich Nelson Tou. RABBIT: An interface for Information Retrieval by reformulation. Ph.D., Massachusetts Institute of Technology , 1982.