

# Challenges and Opportunities of Context-Aware Information Access

Gareth J. F. Jones  
Centre for Digital Video Processing  
Dublin City University  
Glasnevin, Dublin 9, Ireland  
Gareth.Jones@computing.dcu.ie

## Abstract

*Ubiquitous computing environments embedding a wide range of pervasive computing technologies provide a challenging and exciting new domain for information access. Individuals working in these environments are increasingly permanently connected to rich information resources. An appealing opportunity of these environments is the potential to deliver useful information to individuals either from their previous information experiences or external sources. This information should enrich their life experiences or make them more effective in their endeavours. Information access in ubiquitous computing environments can be made "context-aware" by exploiting the wide range context data available describing the environment, the searcher and the information itself. Realising such a vision of reliable, timely and appropriate identification and delivery of information in this way poses numerous challenges. A central theme in achieving context-aware information access is the combination of information retrieval with multiple dimensions of available context data. Potential context data sources, include the user's current task, inputs from environmental and biometric sensors, associated with the user's current context, previous contexts, and document context, which can be exploited using a variety of technologies to create new and exciting possibilities for information access.*

## 1 Introduction

In a few short years information retrieval systems have emerged from being the domain of librarians and information search specialists to become a near ubiquitous tool of all computer users. Current commercial information retrieval systems, such as *Google*, can achieve high accuracy document retrieval from a collection of billions of web pages with remarkable speed in response to short unstructured search requests. However, impressive as current search engines are, they are actually quite restricted in their

functionality. Significantly they take no account of the *context* in which the search is taking place. This means that the same documents are retrieved for every user entering a particular search request, whatever their background may be, whether they have viewed individual documents before, wherever the user may be, and whatever they may be doing or whatever their state of mind may be. Users, of course, live and work in real-world contexts. For example, they may remember where they were when they read something before or what they read just afterwards, their mood may change while carrying a search for information, or the relevance of information may vary depending on the location or time of retrieval. One of the aims of current research in information access is to take account of these and other context factors to improve the accuracy and efficiency of the user's search experience.

The limitations of current search technologies will become increasingly apparent as advances in computing systems extend the scale of and environment for information management systems. While digital information archives such as the World Wide Web are continuing to expand rapidly, ongoing improvements in the capacity of information storage devices, technologies for information capture and distribution networks mean that the already rapid growth in the amount and variety of accessible online information is set to accelerate dramatically in the next few years. At present, this rising data capacity is enabling significant increases in the storage of multimedia information in the form of images, audio and video documents. At the same time, advances in networked computing are creating infrastructures that mean users are effectively becoming embedded in active pervasive computing environments. These embedded users are permanently connected to rich online information resources, but more significantly these environments enable the capture of a user's activities and physical states, attributes of their environment and features associated with documents. Captured data can range from collecting documents that a user accesses or creates on a computer, to information captured from sensors, such as

multimedia video and audio recordings of the user's activities or simple environmental sensors monitoring location, light levels, temperature, etc. Data captured in this way can be used to record context data associated with user activities, the context in which new documents or content are generated, or the context in which existing material is accessed. This context data can potentially be used to identify relevant content in subsequent information retrieval phases. From a complementary perspective context data captured from the user during the search process itself can be used to guide the retrieval process to help identify content relevant to the user in their current context. The full realization of utilizing all the data available from these networks presents many challenges to information access technologies, but the potential rewards from successfully doing so represent a compelling reason to pursue this research.

This paper begins by considering the motivation for and applications of context-aware information access (CAIA). It then explores issues in context-aware information indexing and subsequent implications for context-aware information retrieval. The paper reviews existing research projects, including some of our own work at DCU, which focus directly on technologies pertinent to this area, and concludes with some thoughts about their integration to form ubiquitous CAIA applications.

## 2 Why Context-Aware Information Access?

The desire to find out about new things or to recall facts or experiences from the past is a natural human characteristic. Throughout history peoples have explored their environment to discover new things; more recently, firstly through libraries, then the published media, and currently via the internet, people have sought to satisfy their seemingly inexhaustible desire for information from larger and more heterogeneous data sources. Whether for work or pleasure, people are always in need of information. The ready adoption of the internet and search engines by both specialist and general users provides ample evidence that if information is available, people will access it. But aside from explicit information sources, such as published books and articles, and informally produced web pages, many other sources of information are potentially available for collection into permanent archives derived from ephemeral activities. These may be personal experiences such as social activities, for example watching movies or sports events, or listening to music can be viewed as activities worthy of recording and collecting into information archives, or more formal activities, such as lectures, presentations, debates and discussions or business meetings.

Beyond this, advances in wearable computing mean that it will soon be possible to capture a digital record of all activities carried out in a person's lifetime. This could in-

clude all the documents ever viewed or written by a person, but also video records of places visited and people met, and much more besides. The contents of these archives can be labelled with context data associated with their capture. Issues relating to these emerging archives, commonly referred to as *Human Digital Memories*, are becoming a major research focus [10]. It is a natural component of this paradigm of information capture that context annotation can also be associated with each occasion on which a document is accessed. Humans make extensive use of context when remembering things, so it is a logical development to argue that information access technologies should seek to do the same [2]. Of course, we often do activities merely to enjoy them at the time, but the act of remembering events is itself often pleasurable. Being able to efficiently retrieve them from an archive "memory" of our past experiences is thus a desirable goal.

As an example of the emergence of a this type of technology and associated user behaviour, consider digital photography. Users are increasingly generating personal digital photographic archives. These collections of files are often archived on computers and various other storage media. As an individual's digital photograph collection grows it becomes increasingly difficult to search the set of available images. Search is currently only possible based on filenames and directory structures. On many occasions it would be desirable to search based on persons, objects or places, or possibly events depicted in photographs. However, automated analysis of the semantic contents of images is currently a long way short of reliable annotation of this sort, and it is improbable that anyone except the keenest photographer would manually generate this level of rich annotation. It is though relatively easy to record context information associated with the taking of a photograph, for example the time at which the image was taken and its location, captured using GPS measurements. This data can be used to look up the associated named location at various levels of granularity via a gazetteer, and absolute and relative descriptions of the time, e.g. *yesterday*, *Wednesday morning*, *in January*, *in the Spring* [2]. Using external sources of information details such as the World Wide Web, prevailing weather conditions can be labelled, and even expected light levels, e.g. *hot weather*, *nighttime*, can be labelled automatically. By gathering context data together associated with multiple images together, further metadata can be inferred, such as labelling of groups of images associated with specific events [15]. Once labelled in this way the user can explore the image collection for one of their own photographs using remembered metadata, or use the metadata to search for suitable images from a published collection, e.g. *springtime photographs of cherry blossom near the Imperial Palace in Tokyo*. Within the CDVP at DCU we are developing a system to support user searching of digital

photograph archives using these context labels within the *Mediassist* project. The current demonstration system enables context based searching of an archive of several thousand images utilizing metadata fields of the form described above. This will be extended to searching with a combination of context metadata and automatically indexed visual content features.

The vision of fully functional HDMs goes far beyond this single task gathering of information and context, to the proactive collection of all details of an individual's life. CAIA can be concerned only with retrieval from existing and next-generation HDM personal archives, but such retrieval systems also have enormous potential for applications able to proactively select and deliver new information, the existence of which the user may be unaware, to them that may be of assistance or interest in their current situation. A limited version of this idea of proactive memory augmentation already appears in the information retrieval agents developed using the *Remembrance Agent* described in [19]. Finding information of this type that relates either to a user's current context or to a specific past experience will only be possible if contextual information itself forms part of the information capture and searching process. For example, users may be able to recall where and when they read a document with much greater clarity than the contents of the document itself. Or a search engine may be able to respond differently depending on how the user is feeling based on biometric responses.

Appropriate selected information will be *relevant* to the user's *current* context. Thus the user may be using a conventional desktop computer and fully engaged in the searching process, but if they are away from their desktop they are probably more likely to be engaged in other activities. However, as discussed earlier, they are still likely to be connected to a computer network via a mobile device. Selection and delivery of relevant information to mobile users carrying mobile computers provides both challenges in terms of appropriate form and timing of delivery, and opportunities to exploit physical context information associated with the user embedded in the sensor rich networked environment.

### 3 Paradigms for Context-Aware Information Access

It should be clear from the discussion so far that there is a complex relationship between context and documents and users. Context is associated with the creation of documents, the contents of documents, relates to previous and current access to documents, the user's previous context, the user's current context and possibly a prediction of the user's future context [4]. In terms of using and interacting with information and associated context, the user's search context is a vital feature. A user working with a desktop computer

will most likely be fully engaged with the search process, and have access to a large high quality screen for data entry, manipulation and browsing. A user accessing information using a mobile pervasive computing device is more likely to be primarily engaged in other activities, and only has access to a physically small device to interact with. In order to better understand these search modalities and begin to establish how permanently available proactive information applications might work, this section explores successively more complex context-aware search paradigms.

#### 3.1 Interactive and Proactive Retrieval of Previously Seen Information

Current search engines make no distinction between users looking for documents that they have not seen before and items that they have viewed previously, possibly many times. Research suggests that on average around two-thirds of information that people utilize is contained in sources with which they have interacted previously [8]. Given that users will often remember context features associated with this information, it makes sense to gather together previously seen information and enable users to search this subset of the information available to them.

This scenario is addressed directly in the "Stuff I've Seen (SIS)" system developed by Microsoft [8]. SIS contains a local index of documents, including emails, web pages, images and diary entries, written or accessed by an individual user. Along with the document contents, SIS stores context metadata of dates, authors, text previews and image thumbnails. Utilization of a local index in this way enables very rapid searching and query refinement. Users are able to restrict searching to individual data sources, and to use standard text information retrieval methods to order lists by document contents, but also to sort the rank and filter documents based on other metadata features. The interface is designed to enable rapid browsing of retrieved subsets of the collection by showing thumbnails of images and short 300 character introductory previews of text documents. By launching local searching whenever metadata filtering widgets are manipulated queries can be efficiently iteratively adjusted with rapid feedback to find desired information.

The functionality of SIS systems is clearly important for HDMs technologies. However, an important feature of the SLS prototypes is the requirement for the users to actively interact with the system to locate the desired information. This means that the user must give their full attention to the search process, and also significantly they must have a large high quality interface to the system available and that sufficient computing power must be available for very rapid iterative searching as the query is refined. While it is likely to be inevitable that the user will sometimes need to work hard with any system to locate more difficult to find

previously seen information, the requirement of this to be the expected mode of searching will limit the scope for the exploitation of this type of system. This will be particularly true for more complex HDMs, which will often include hard to browse multimedia content, and need to be available to mobile users.

An alternative to interactive searching of previously seen “stuff” is provided by the Remembrance Agent (RA) applications developed by Rhodes [19]. The basic search paradigm of RAs is to look at the content and other data associated with the user’s current working context and proactively search for available information which might currently be of interest to the user. The basic version of the RM was implemented into the `emacs` text editor. While a user is working on a new document, the RM takes the text that is currently being entered and uses it to form a search request for documents relating to this content which may be of interest to the writer. For example, this would search for related references while developing a new academic paper. An development of the basic RM is seen in the “Marginal Notes” system which augments a standard web browser. When a web page is downloaded it is rewritten before it is displayed to include a margin at the side of the window. The web page is segmented based on the HTML markup and proactively formed into “requests” and used to search available content. Notes relating to the highest scoring documents for each request are then inserted into the margin. The note contains a brief description of the suggested text, a series of circles representing the relevance of the suggestion, and a link to obtain more information. The note consists of a subject, date, and author for the suggested text, the exact composition of the note can be customized. Placing the mouse over a suggestion note produces a list of the five keywords that were most important in determining its selection. Clicking on a suggestion note opens a new browser window that displays the full text of the associated document.

While these RM systems introduce the concept of proactive searching of previously accessed content, they are still implemented on desktop computers and do not take account of context information associated with the documents or the user’s current context.

### 3.2 Mobile Users

As well as the established paradigm of using search engines on desktop computer, as described earlier, mobile users are increasingly permanently connected to ubiquitous wireless computer networks via their mobile phone or PDA. This gives them direct access to information that may be of assistance to them in their current context. Working with these small devices means that high levels of interaction are difficult, and it is thus important to exploit whatever information is available to maximise retrieval precision. Since

the information that a user is looking for when mobile will often relate in some way to their current context or an anticipated future one, it is natural to seek methods by which this data might be incorporated in information searching.

Proactive retrieval, where the user does not have to break off from their current activities to look for potentially useful information, is thus a particularly appealing scenario for mobile information delivery. This scenario builds on the basic Remembrance Agent methods outlined in the previous section to incorporate context data. A mobile version of the RM incorporating context is provided by the “Jimminy” system. Jimminy searches for information based on a person’s physical environment, including location, people nearby, time of day, and subject of conversation. Processing is performed on a shoulder-mounted “wearable computer” and suggestions presented on a head-mounted display. The ultimate goal is that all information about the wearer’s physical environment will be available to Jimminy through automatic sensors. Jimminy is a general architecture able to plug-in any sensor that can be attached to a wearable computer. Information not available from sensors can be entered into the system by hand, for example the topic of conversation can be entered manually in the form of text notes. Notes entered into the wearable computer can be manually tagged with people present, subject, location, and time stamped. Search is performed using a retrieval engine called *Savant* which incorporates context fields in the retrieval process.

Jimminy provides a starting point to the potential of mobile proactive memory augmentation and information searching applications. However, it falls short of addressing many significant issues. In our previous work we have explored a number of aspects of context and information retrieval for mobile applications [4][13]. One of the important conclusions of Rhodes’s work was the observation that for mobile users *precision is generally more important than recall*. This will be especially true of ubiquitous proactive retrieval where the user has not actively entered a search request. In this scenario we regard each retrieval brought to the attention of the user as an intrusion. Therefore the aim should be only to deliver items that are both relevant and cannot be handled automatically via some form of agent acting on the user’s behalf. Further, even when it is determined that the information needs to be delivered to the user, it still needs to be presented in the most appropriate mode [13].

While it has been demonstrated within *Mediassist* and elsewhere that context can be used effectively in information searching, it has not been formally demonstrated that incorporating context into the search process is able to improve precision in ranked retrieval, although we hypothesize that it is. There are a number of open research questions here. Which context fields should be incorporated into the retrieval process? How should this data be integrated into retrieval models? What other issues need to be addressed

for effective context-aware retrieval on mobile devices?

### 3.3 Human Digital Memories

Beyond retrieval in the context of previously seen documents and the issues of retrieval to mobile devices, this section provides an introduction to the search needs of *Human Digital Memory (HDM)* technologies. HDM systems are typically envisaged as creating archives of every document created or viewed by a user during their lifetime together with temporally related documents, and increasingly include multimedia video, image and audio content captured from the user's personal life experiences. This content is captured in real-time during the user's daily activities. Life-time personal archives created in this way are going to be, by today's standards, physically huge; but the greater long-term challenge is automatically finding interesting or relevant content from within these large highly heterogeneous collections, and delivering it in a form that is useful.

The emergence of these technologies realizes the vision of Vannevar Bush's *Memex* described in 1945 as "a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility" [5] Bush could not foresee exactly how the Memex would be achieved, but he did envisage two fundamental components: *annotation* and *links*. Annotation makes the archive searchable and describes individual items and potentially collections of items. Linking items within an archive provides a vital source of context information which will often make the stored items meaningful. Examples of current research exploring HDM systems include the Microsoft MyLifeBits project [9] and personal Life Logs at the University of Tokyo [2]. These projects are raising many interesting questions relating to the reuse of captured data, but research into useful applications exploiting these resources is currently at a very early stage.

A question which often arises with respect to HDM archives is: do they have any real value? The best answer to this probably lies in recognising the value that individuals already place on photo albums, framed photos, personal letters, newspaper clippings, and home movies and videos. While many of these items are accessed only rarely, perhaps only a few times in a lifetime, they are treasured [9].

**MyLifeBits [9]** The MyLifeBits project is an effort to implement a personal digital store of all information a user encounters in their lifetime. Assuming that the requirements of sufficient storage capacity have been met, the next requirement is to consider how to locate interesting items within the archive of all this information. It is daunting to attempt any type of organisation. Given the scale of the archive it is not only difficult to find desired items, and in-

deed the existence of many will have been forgotten about. The user might even have difficulty recalling the significance of individual items when they are retrieved from the archive. Given these considerations, MyLifeBits is designed based on a number of guiding principles:

- the user must not be constrained to organize items within a strict hierarchy.
- there should be many ways of visualizing the contents to help understand the significance of items.
- the value of non-text media is dependent on annotations.
- authoring tools create two-way links to media that they include in new media - referred to as "transclusion".

Searching for and realizing the true value of items within large collections will depend on the quality of these annotations, but also vitally, as realised by Vannevar Bush, on the links between items. A photograph may be labelled with the date on which it was taken, which may remind you of the season, or an event. However, by linking other photos and items associated with this event, and possibly related events involving the same people, a "story" can be constructed out of the items. These links are crucial because they can provide context and commentary for individual items. Pointers from video keyframes to the original video provide the full context in which it was captured. For others, perhaps children or grandchildren, in the future to have any appreciation of the contents of a parent's HDM, annotations and stories are vital. Additionally, sharing of content among those involved in the items in the collection, colleagues, family and friends, can result in multiple annotations and lead to snowballing effect on the value of the content.

Links between items are fundamental to the MyLifeBits database. A link indicates that one item annotates another one. One item can annotate any number of other items, and may be annotated by any number of others. Searching on the items is currently limited to allowing multiple views based on timelines and searching based on text annotations. The system encourages manual text annotation of items. It also supports audio annotations which are passed to a speech recognition system to make them searchable.

**Life Logs [2]** Unlike the MyLifeBits project which is exploring the development of complete HDM archives, the University of Tokyo Life Logs focuses on the capture and searching of personal life experience multimedia collections recording daily life. A prototype system captures data from a wearable camera, a microphone and various context sensors, including location via a GPS receiver, an acceleration sensor, a gyro sensor and a brain wave analyzer. The user is also able to add text annotations to data

as it is captured. Data is recorded onto a laptop computer which also harvests further information from the internet including, prevailing weather conditions at the user's location, news from the day, all web pages browsed by the user and other online news sources. The system can also record copies of all files opened and edited by the user including emails and word processed documents. GPS locations are converted to textual descriptions using gazetteers to identify regions, but also details of buildings and businesses. Similarly, times can be converted from specific dates, hours, minutes and seconds, and to relative days, weeks, months and seasons. Since indexing and interpreting much of the captured video and audio media is difficult, computationally expensive and often highly subjective, one of the effective implemented features of this system is to search using the context data. For example, they would be able to look for an item recording *a spring meeting on a sunny day outside an Italian restaurant in Ginza*.

One of the key issues addressed in the life logs research is working with the sheer volume of captured multimedia data [1]. While use of the captured context can lead to useful retrieval based on remembered cues, the amount of content available for browsing is vast. Stated bluntly, manual searching of a one month video would typically take longer than the length of the recording. This problem becomes even more apparent when considering that the archive might include years of data, and possibly be gathered from many users. It is thus desirable to identify key material from the archive based on some form of summarization criteria. Objective analysis of the content is unlikely to be able to distinguish between scenes that the user regards as significant or insignificant. The key idea introduced in [1] is the desire to identify key scenes from the life log based on the user's subjective experience. This is based on collected brain wave activity which is interpreted as indicating whether the user is paying attention to the current scene around them. The system proceeds by first using low-level feature analysis to identify likely shot boundaries in the record video data. The brain wave activity level is then used to identify shots likely to be the most significant to the user. Various methods of content selection are explored in the original reported work.

In more recent work the project explored identification of keyframes within video based on spatio-temporal sampling of context data [2]. Various criteria for selecting keyframes including absolute changes in time and location, and rates of change. Once identified keyframes can be subjected to computationally intensive processing to extract significant visual features. One of the features currently under investigation is the identification of conversation scenes. The extent of the visual features that can currently be reliably identified is obviously limited, but the general framework developed here provides an excellent testbed for exploring future developments in image processing.

## 4 Further Sources of Context

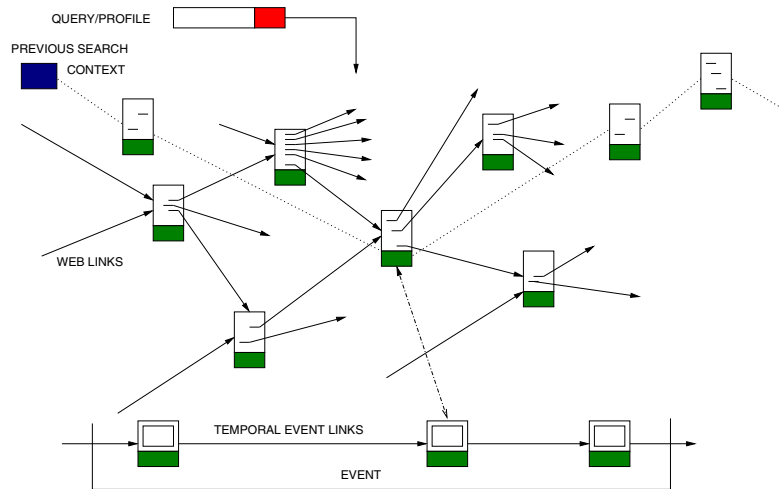
The context features discussed so far are generally easily identified and are becoming increasingly easy to capture. However, developments in sensor technologies and further exploration of context associated with documents and searching can be used to gather and annotate a range of additional context attributes which might usefully be exploited in retrieval.

### 4.1 Biometric Context

Recent studies have demonstrated that a number of classes of a person's emotional states can be recognised automatically using non-intrusive physiological measurements [14][16][18]. e.g. recognising surprise or amusement. These generally focus on easily measurable physiological features, and, unlike the work described in [1], are largely non-intrusive. In a previous experimental study we illustrated that cognitive engagement in an emotional state can have a detrimental impact on the ability of subjects to successfully undertake a decision-making task which required them learn about the task environment [12]. This task was based on a earlier study used to explore the decision making ability of affect-impaired patients [7]. This earlier study showed a clear correlation between galvanic skin conductivity and the anticipated outcome of actions in the task for normal affect-unimpaired subjects. Taking this existing work together, we hypothesize that users of information retrieval systems produce measurable physiological responses as their search progresses. Utilizing this data would require the system to differentiate between signals arising from the search process itself and the impact on the user of the content of retrieved information. The physiological measurements could then be correlated with a user's emotional states arising from the search process, e.g. success, failure, frustration, etc.

Psychological studies have shown that emotional states can be interpreted in terms of three dimensions: arousal, valence and control [3]. We are currently exploring the interpretation of audio signals using these dimensions for affect labelling of multimedia data [6]. At present, we are working with movies and TV footage, but this system could be adapted to process the contents of HDM archives. Other work already incorporates visual features in its interpretation of arousal and valence [11], and we intend to extend our system to include these features as well. Labelling of documents in this way would enable searching for retrieval of items captured with a strong emotional component, e.g. excitement or sadness.

Once the physiological responses to events during interactive information searching have been understood, it is interesting to consider how search engine algorithms might be made adaptive to respond to this information to improve



**Figure 1. Multiple link and context sources in a context annotated document collection.**

the effectiveness of search. For example, to reduce anger or frustration, possible methods available include emphasizing the presentation of completely new information which is highly ranked rather than presenting “more of the same” information which might be also be highly ranked.

#### 4.2 Context of Information Access

Context can be associated with documents when they are created, and with the user when a search request is entered. An extension of context labelling is to dynamically annotate documents with details of the context in which they are retrieved. Exactly what annotations should be made would obviously depend on the data available, but we could think in terms of the standard attributes of location and time, and also biometric data gathered from the user. Biometric data could be used both to gauge the likely relevance of the retrieved items, and to record the user’s reaction to the retrieved information itself. An additional feature could also be to annotate the document with the query used to retrieve it, and whether it was found to be relevant to this query. Note, this does not imply that this document should always be expected to have the same relevance relationship to this query, but it is potentially useful information.

Recent years have seen much research into the exploitation of inter-document links to improve information retrieval effectiveness. The most well known method of using linkage structures is the *PageRank* algorithm [17] which formed part of the original *Google* search engine. The use of links between items in a HDM is emphasised in *MyLifeBits* [9] to describe stories, and the linking of photographs within digital collections is a current active research area, for example [15] and within our own *Medias-*

*assist* project. In this work inter-document links are concerned either with links generated when items are created, logical links between items forming an “event” or “story”, and links created by web users to point to significant or related pages. Related retrieved items could be dynamically linked to other documents retrieved in response to a query. In effect the query places documents into a context, preserving this context by linking retrieved documents, particularly if they are labelled as relevant, enables richer descriptions of stories and events to be developed, and potentially sharing of annotations. Figure 1 illustrates this situation of multiple types of inter-document links and context annotations. The link structure will become ever richer as documents continue to be retrieved in response to further search requests. As discussed earlier, when trying to recall the details of a specific document, a user will often recall the context in which it was viewed, by annotating documents with related context data each time they are accessed, this data will be available for use in future search requests when the user wants to find the same document, or contextually related (encompassing all possible interpretations of context) documents, again in the future. The overall aim must be to more reliably retrieve documents in future search operations.

### 5 Ubiquitous Context-Aware Information Access

Proactive context-aware delivery of information to mobile users from HDMs and other sources combines all the issues explored in this paper. A non-intrusive version of this system would form a type of augmented memory or information assistant.

Putting together a system which can deliver information to a user is only part of the requirement; the information delivered must be useful and enrich the user's experience of their environment. Information delivery must be reliable and timely, and in a form suitable for the user's current activities. This raises issues of detecting the user's likely current activity, and deciding if or when to deliver information, and how much information and in what media it should be delivered. We thus arrive at a system incorporating some form of personal assistant agent which can manage information searching, selection, and delivery, and determine whether to summarize the information, or transform it into another media. For example to present details in a map as a sequence of audio directions to a user who cannot safely or conveniently look at a map while driving or walking, or as a textual list to a navigator. These issues are explored in more detail in our earlier work [13].

Work on HDM searching currently makes much use of interactive exploration of context as part of the query development process [2][9]. In order for these archives to be useful to users in their daily life, retrieval from HDMs must be effective when using mobile devices with the attendant difficulties of user interaction explored earlier.

Work carried out to date has only begun to explore the potential of HDM technologies, information access on mobile devices and the exploitation of context in information access. There are many interesting areas for further exploration. For example, a speculative idea for intimate integration of biometric data into the search process arises from [7]. In this work, Antonio Damasio speculates that humans re-experience a weaker physiological manifestation of the states associated with something that they are remembering. An interesting question raised by this observation is whether this experience might be captured during search as the user remembers a document what they are looking for and then be exploited to help them find it again.

**Acknowledgement** This work is partially funded by Science Foundation Ireland under grant 03/IN.3/I361.

## References

- [1] K. Aizawa, K. Ishijima, and M. Shiina. Summarizing Wearable Video. In *Proceedings of ICIP 2001*, pages 398–401, 2001.
- [2] K. Aizawa, D. Tancharoen, S. Kawasaki, and T. Yamasaki. Efficient retrieval of life log based on context and content. In *CARPE'04: Proceedings of the the 1st ACM workshop on Continuous archival and retrieval of personal experiences*, pages 22–31. ACM Press, 2004.
- [3] M. M. Bradley. *The Emotions: Essays on Emotion Theory*, chapter Emotional Memory: A Dimensional Analysis, pages 97–134. 1994.
- [4] P. J. Brown and G. J. F. Jones. Exploiting Contextual Change in Context-Aware Retrieval. In *Proceedings of the 17th ACM Symposium on Applied Computing (SAC 2002)*, pages 650–656. ACM, 2002.
- [5] V. Bush. As We May Think. *The Atlantic Monthly*, 176(1):101–108, 1945.
- [6] C. H. Chan and G. J. F. Jones. Annotation of Multimedia Audio Data with Affective Labels for Information Management. In *Proceedings of the Fifth International Workshop on Pattern Recognition in Information Systems (PRIS 2005)*, 2005.
- [7] A. Damasio. *Descartes' Error*. New York: Grosset-Putman.
- [8] S. Dumais, E. Cutrell, J. Cadiz, G. Jancke, R. Sarin, and D. C. Robbins. Stuff I've Seen: a system for personal information retrieval and re-use. In *SIGIR '03: Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval*, pages 72–79. ACM Press, 2003.
- [9] J. Gemmell, G. Bell, R. Lueder, S. Drucker, and C. Wong. Mylifebits: fulfilling the Memex vision. In *ACM Multimedia '02: Proceedings of the Tenth ACM International Conference on Multimedia*, pages 235–238. ACM Press, 2002.
- [10] J. Gemmell and H. Sundaram, editors. *CARPE 2004 The First ACM Workshop on Continuous Archival and Retrieval of Personal Experiences*. ACM, 2004.
- [11] A. Hanjalic and L.-Q. Xu. Affective Video Content Representation and Modeling. *IEEE Transactions on Multimedia*, 7(1):143–154, 2005.
- [12] R. A. Hase, G. J. F. Jones, and J. J. Loose. Exploring the Impact of Emotional State on Sensitivity to Future Consequences. In *Proceedings of the Symposium on Emotion, Cognition and Affective Computing, at the AISB'01 Convention*, pages 98–104.
- [13] G. J. F. Jones and P. J. Brown. Context-aware retrieval for ubiquitous computing environments. In *Workshop on Mobile and Ubiquitous Information Access, Mobile HCI 2003 - 5th International Symposium on Human Computer Interaction with Mobile Devices and Services*, pages 227–243. Springer, 2003.
- [14] K. H. Kim, S. W. Bang, and S. R. Kim. Emotion Recognition System using Short-Term Monitoring of Physiological Signals. *Medical & Biological Engineering & Computing*, 42:419–427, 2004.
- [15] M. Naaman, Y. J. Song, A. Paepcke, and H. Garcia-Molina. Automatic Organization for Digital Photographs with Geographic Coordinates. In *JCDL '04: Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries*, pages 53–62. ACM Press, 2004.
- [16] F. Nasoz, K. Alvarez, C. L. Lisetti, and N. Finkelstein. Emotion Recognition from Physiological Signals for Presence Technologies. *International Journal of Cognition, Technology, and Work - Special Issue on Presence*, 6(1), 2003.
- [17] L. Page, S. Brin, R. Motwani, and T. Winograd. The Pagerank Citation Ranking: Bringing Order to the Web. Technical report, 1998.
- [18] R. W. Picard, E. Vyzas, and J. Healey. Toward Machine Emotional Intelligence: Analysis of Affective Physiological State. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(10):1175–1191, 2001.
- [19] B. J. Rhodes and P. Maes. Just-in-Time Information Retrieval Agents. *IBM Systems Journal*, 39(3&4):685–704, 2000.