Using context awareness to enhance visitor engagement in a gallery space

Peter Lonsdale[1], Russell Beale[2], Will Byrne[1]

[1]School of Engineering and [2]School of Computer Science, University of Birmingham, Edgbaston, Birmingham B15 2TT
Email: *{p.lonsdale@, r.beale@cs., w.f.byrne@}bham.ac.uk*

Context-awareness can greatly enhance the usability of mobile devices by making it possible for users to continue with other activities without having to pay too much attention to the device. At the same time context-aware applications can provide timely support for user activities by responding to changes in the user's state and acting accordingly. We describe our work on developing a generic context awareness architecture that is being deployed in a gallery space to enhance learner engagement with the gallery exhibits. Our system makes use of contextual information to determine what content should be displayed on the device. Users can also navigate this content by explicitly changing their context in the dimensions of physical location and dwell time. Visitors have the opportunity to physically interact with the abstract information layer that is overlaid on the gallery space. The system also actively encourages movement in the gallery by identifying links between paintings. We describe our architecture, implementation, and the design challenges faced in deploying this system within a gallery.

Keywords: context awareness, mobile learning, museum, mobile usability, PDAs, location awareness.

1 Introduction

Context awareness is a relatively nascent field of research that centres on the use of information pertaining to the user and their environment to drive the behaviour of a device or system (for reviews of context-awareness applications and research perspectives, see Chen and Kotz [2000] and Dourish [2004]. By taking account of contextual information, systems can be made easier to use, and can provide more appropriate responses, than if they respond passively to user requests. There is currently considerable interest in the use of context awareness to provide enhanced usability in this way, and in particular developers of mobile applications are looking to context awareness to provide solutions to the specific problems posed by the design limitations of mobile computing devices.

Context awareness is especially important for mobile devices because it allows us to overcome the usability problems associated with small, handheld devices[Sharples, M. and Beale, R. (2003) *Journal of Computer Assisted Learning,* 19, 392-395.], and also to make effective use of the user's physical and social surroundings to provide timely support to their activities. We can provide enhanced user support by using contextual information to drive mobile applications, and simultaneously we can exploit physical context as a means to interact with the application. For example, a user's location can drive delivery of content on a device, and at the same the user can utilize the system's sensitivity to their location as a navigation tool, selecting different items of content by altering their physical location.

1.1 Context awareness for mobile learning

Mobile learning is too often conceived of as simply the mobile equivalent of e-learning. The assumption is that learning can be delivered through content displayed on mobile devices in the same way as it is displayed on other systems such as desktop PCs. However, the use of mobile computing devices is qualitatively different to the use of other computing devices, and we must take account of this when developing m-learning applications. We must also consider that a user with a mobile device is often much more influenced by their surroundings than a user of a desktop PC might be.

Mobile devices such as phones and PDAs are used in a huge variety of settings and environments, and we cannot rely on having the user's full attention. Mobile learning is not something that can be delivered, it is something that might happen, given the right combination of learner, surroundings, content, and activity. This serendipitous nature of mobile learning is further enforced by the very informal way in which people use things like mobile phones and PDAs. They are not setting out to learn something, they are often engaged in something else entirely, and we must make the best use we can of the devices they have to hand to support their activities [Rogers ref]. We can do this using context awareness to ensure that the device is always ready with relevant information, but does not need to distract the user in order to achieve this readiness.

The EU IST Project MOBIlearn [Bormida ref] has focused on the development of a large scale platform for delivering learning content to learners with mobile devices, including mobile phones, tablet PCs, and PDAs. As part of this work, the University of Birmingham has developed an architecture for context awareness that is currently being deployed at the Nottingham Castle Museum gallery to provide an enhanced and more engaging experience for visitors to the gallery.

Our aims have been:

1. To provide timely support to the user
2. To allow the user to maintain their attention on the world
3. To allow the user to inspect, understand, and alter the current context model for their own purposes

We describe our approach to implementing a generic context awareness architecture for mobile learning applications, and discuss its use in deploying a system to enhance visitor experiences at Nottingham Castle Museum gallery.

1.2 Context awareness for museum and galleries

Museums, galleries, and heritage sites seek engage visitors in the artefacts they exhibit, as well as encourage participation in the learning space provided. The use of computer technologies in museums is not a new concept, and kiosk-based content presentation and interactive exhibits are a common sight. Mobile devices offer opportunities to provide technological means to engage visitors whilst they are situated within the gallery space itself. Moreover, the small size and portability of these devices means that we can seek to engage the user without distracting their attention from the exhibits they are trying to enjoy.

Context-aware mobile applications have been used effectively to deliver supporting information to tourists in the form of location-aware tourist guides, for example the CyberGuide project [Abowd et al., 1997]. Similarly, location-aware applications have been used to deliver content that is appropriate to a visitor's particular location within a gallery space – notable examples include Tate Modern's multimedia pilot study [Proctor & Burton, 2003], the CAERUS system at the Winterbourne Botanic Gardens in Birmingham [Naismith & Smith, 2004]

Wearable and mobile computers have also been used to provide augmented experiences that go beyond basic location awareness. Baber et al [2001] describe a system that combines location-dependent content delivery with profiling of visitor needs, to provide a visit tailored to individual requirements. Oppermann et al [1999] describe a system that uses contextual information as the basis for supporting the user but not distracting them, whilst MacColl et al [MacColl 2002] describe their experiences of combining context and virtual presence.

Location remains the primary feature of context that is exploited in most context-aware applications[dix 2000, chen 2000, selker 2000, Bristow 2002]. To provide effective support for visitors to museums and galleries it is crucial to know where they are. Knowing what area a visitor is in means we can offer appropriate content and suggest possible activities. Knowing exactly which artefact a visitor is currently looking at means we can offer content and activities specifically for that artefact.

1.3 Uses of context

A visit to a museum is not just a series of stops in front of artefacts. The experience has a beginning, a middle, and an end. It is a process. We have sought to address this by considering visitor movement within the gallery. We began by designing our system to use location and timing information to provide appropriate content to users with mobile devices in a gallery space. What became clear was that the delivery of content in this way allowed us to encourage visitors to interact with the artefacts in a different way. A review of the content supplied to us by the gallery indicated that many of the paintings on display shared interesting histories or were linked in some way that was never made visible to visitors. By flagging these connections to the visitors we are able to encourage greater movement between the paintings on display, beyond the basic linear path that most people follow in the gallery. This physical engagement with the learning space is an often neglected facet of learning. [??Mumford 2003?]

In our preliminary trials of our prototype, we discovered that because the system relies on context to deliver content, changes in that context can be exploited by the user themselves to deliberately trigger content changes and hence move to another item of content when desired. In other words, context becomes not just a mechanism for the system to select content, but also a tool with which the user themselves can navigate the information space. These two aspects of context aware gallery exploration, *context as content selector* and *context as navigation tool*, have driven our subsequent development of a system to support visitors to the gallery. The concept is that an information space that is overlaid on the existing physical space of the gallery can be navigated through physical means, engendered by the implementation of a context-sensitive application. Mobile devices are hard to use, because they have small screens, and the user is usually trying to do something else at the same time as navigate the onscreen menus. By using physical movement as an interaction method, we can give people a new way to interact with the information space that we have overlaid onto the physical gallery space.

2 Context Awareness Architecture

Context aware applications typically involve the use of rulesets or some other kind of matching system to generate appropriate system *responses* (e.g. content display or option selection) from appropriate *stimuli* (e.g. changes in location, orientation, lighting levels, user input etc). This approach requires the definition of fairly rigid rules (or their equivalent) and an exhaustive set of possible responses that the system can make. Our approach has been driven by the need to support the process of learners moving through a learning space, and hence our model and architecture for context awareness is much more process centred.

Another motivating factor in moving away from programmatically defined rules is the desire to support content developers and experience designers who wish to make use of context aware applications without wanting to engage in software development.

We have devised an architecture (described in detail in [Beale & Lonsdale, 2004]) for context awareness that involves the definition (in textual form) of a set of software objects called Context Feature Objects. Each of these Context Feature Objects responds to a specific stimulus from actual context data, and responds by searching for matching metadata tags on the available content that the system can currently deliver. Any match results in the current score for that item of content or action being increased. When all available content and actions have been scored, those with higher scores are deemed to be of more relevance to the current context than those with lower scores. This scoring or ranking process occurs every time the context changes, e.g. whenever the user moves to another location.

The context awareness system is configured by specifying a set of context feature objects and link objects using a structured syntax that the system parses at runtime to generate actual software objects that perform the context awareness processing. In specifying context feature objects, it is necessary only to know the name of the metadata tag that is appropriate and the range of values that a context feature object should respond to. Links between context feature objects are similarly defined. In this way we have provided a non-programmatic interface to the context awareness system, and one that could easily be translated into even more usable tools such a graphical user interface for the configuring of context aware applications.

Contextual data itself is assumed to be gathered by separate systems, and is input into the context awareness architecture in a generic fashion by specifying simply a name for the data and then supplying its value. Context Feature Objects that are able to respond to the type of information passed in will do so. In this way the context awareness architecture is only loosely coupled to the technical infrastructure which provides the actual context data, and different sources of context data may easily be substituted at any time. From a technical perspective, this functionality is further enhanced by the use of a web services architecture to deploy the system. This means that communication with the context awareness architecture is easily achieved through standard protocols and data formats.

2.1 Conceptual context model

The easiest way to understand the contextual approach we have taken is through the metaphor of a movie. The movie itself has a main theme, and a variety of subplots and threads running through it. This is equivalent to the overall context. It is dynamic, changing over time and with the interactions of the participants, where history is important. A scene in the movie corresponds to a *context state*: a specific set of themes and characters are to the fore and have primary importance. A scene from the movie has these key characters in it, plus some props – this corresponds to what we call the *context substate*. Thus, as in a movie, the whole movie is needed for a full understanding, but a lot of information does exist in a single frame.

2.2 Context Feature Objects

Our software architecture comprises a set of software objects called context feature objects (CFOs) that correspond to real-world context features relating to the learner's setting, activity, device capabilities and so on to derive a context substate, as described above. Data can be acquired through either automated means (for example sensors or other software subsystems) or can be input directly by the user.

This context substate is used to perform first exclusion of any unsuitable content (for example high-resolution web pages that cannot be displayed on a PDA) and then ranking of the remaining content to determine the best *n* options. This ranked set of options is then output to the content delivery subsystem.

2.2.1 Types of context features

Context feature objects are either *excluders* or *rankers*. Items of content that are deemed entirely inappropriate for the current context are excluded. That is to say they are removed from the list of recommended content and not subject to any further consideration – items that match a single exclusion criterion will not receive any further rankings and will not be recommended no matter how high a score they receive, and so exclusion is qualitatively different to simply receiving a low or zero ranking. Content remaining in the list after the exclusion process is then ranked according to how well it matches the current context. The ranking process simply increments the score of each item of content that has metadata matching the stimulus values of any particular context feature. The size of the increment depends on the *salience value* of the context feature doing the ranking. Individual CFOs can have their salience values changed so that they exert more influence on the ranking process. Any individual CFO can be de-activated at any time so that it has no effect on the exclusion or ranking processes.

A CFO has a set of possible values, and an indicator of which value is currently selected. It is also possible for CFOs to have multiple sets of possible values, with the current active set being determined by the current value of another linked context feature. Whilst this has no bearing on the recommendation process, it is important in terms of providing an inspectable model of the context state to the user, who can observe the influence of one context feature on another. For example, options relating to current activity can change depending on the user's current location.

2.2.2 Linked context features

Each context feature object responds to only one metadata tag and performs either an exclusion or ranking function. To achieve more complex filtering of content, CFOs can be linked together so that their function can depend on the state of other context feature objects. Link objects are used to send either the *values* of context features or the *time* they have held that value to other context features. Criteria on that link determine whether action should be taken.

For example, we might have a context feature that responds directly to input from a sensor network specifying the location of the user. Another context feature infers the level of interest of the user by taking input from a link that acts on the *time* the location feature has had its current value. A user dwelling in one place for a longer period implies a higher level of interest in that location. A third context feature may respond to user input that can over-ride the inferred level of interest – this uses a link object that acts on the *value* input by the user. Conflicts between links and context features are resolved using salience values which specify the relative importance of each. These salience values are at present specified by the designer(s) of the context-aware experience, but more automated methods of conflict resolution could be employed in future iterations.

2.4 Output

The ordered list of ranked items of content is passed to delivery subsystems for use in determining exactly what content should be made available to the user. In this way, the context-awareness sub-system has no way of specifying exactly what is made available – the system is intended only to make recommendations to the system and to the user. This method of recommendation is preferred so that should the system make a mistake, and make inappropriate recommendations, its output does not override selections made elsewhere in the system (for example, the user might specify a particular page of content and then not want that item to be replaced by another).

It should be clearly understood that the recommendations made are not only done on content – recommendations can also determine new navigational strategies through the virtual or real space. We are not concerned with only filtering content, but in the more general question of providing appropriate support, which may be re-ordering information, offering it in a different order, or directing the user to another part of the physical space – which will in turn affect the context system [lonsdale et al 2004??].

3 In the gallery

We have deployed our context awareness architecture in the gallery space at Nottingham Castle. Our intention is to provide visitors to the gallery with an enhanced experience through the utilisation of contextual information to drive the behaviour of their mobile device. Content and options displayed on the screen will be tailored according to the user's current context, and users are also able to make explicit use of the context sensitivity to drive the behaviour of the device themselves.

3.1 Designing the experience

We have consulted with the curators at Nottingham Castle Museum to ensure that our system will deliver appropriate support to visitors in the gallery space. Several issues arose during our consultation, of which two are immediately relevant to the design of the context aware visitor experience:

1. *Lack of focused attention:* visitors will usually enter the gallery space via one door, move through the space in a linear way, and then exit without really paying much attention to what they see on the way
2. *Deadspots:* certain artefacts within the gallery are often overlooked by visitors, for a variety of reasons; positioning, lighting, or other factors.

We wanted initially to use our context awareness system to attempt to overcome these issues, and thus provide a more engaging experience for the visitors whilst at the same time addressing these areas that concern the gallery staff.

A crucial part of the design centres on the fact that visitors move through a physical space. This movement was determined to be the primary context feature for our system to use. In addition, movement itself is not constrained to two or even three dimensions – visitors' movements can be described also in terms of the fourth dimension, time. The particular path a visitor takes through the gallery, the time they spend at individual paintings, and whether or not they retrace their steps can all be used to drive a context aware application. Our system has been set-up to deliver appropriate content using the following principles:

1. *Which painting is the user currently closest to?* This is determined from our positioning system as described below. The system is able to provide accurate data about which painting the user is currently closest to.
2. *How long has the user been in their current position?* An increased dwell time at a specific painting is assumed to indicate a higher level of interest in that painting.
3. *Has the user been in this position before?* If the user has been to a painting before, the content they viewed on their previous visit can be used to determine the appropriate level of content to display this time. Previous content can also be offered for review.

As well as using context awareness to determine what the device does, we are also exploring the use of context awareness as a means to physically engage the learner in the learning space, and to encourage movement within that space.

3.1.1 Encouraging movement within the space

Mobile devices are often deployed in museums and similar locations as a means to deliver content or provide some other element of interactivity to the exhibits. But delivering content means that we are in danger of replacing hands on interaction with 'heads down one-way transmission of information' [Hsi, 2003]. Instead, what we can do is to use the device and the content it can display to cause the visitor to see the artefacts in a different way, and to expose the links between paintings that were not visible without the use of the technology to point them out.

This functionality has been implemented through structuring the audio content provided by the device to highlight links with other paintings in the gallery space. Users are expected to navigate to the other paintings without additional assistance, which is in part the reason we have seen the use of the context sensitivity as a navigation tool.

3.1.2 Enabling navigation through physical movement

We have observed that users wanted to navigate the information space by physically changing their context so that they were effectively driving the system through physical actions. Our application already supported this through being sensitive to context changes, but to further enhance usability in this area we have explored the use of salient contextual information on the user interface so that users can monitor the state of the context system and determine whether they have achieved the state they are aiming for. In this case, it became necessary to indicate to the user the exact location the device was currently registering for them, whether it thought they were moving or stationary, and how long it thought they had been in that location.

3.2 Deploying the experience

To provide the functionality described above within our context awareness architecture, it is necessary to define two Context Feature Objects to monitor Painting and Interest. The Painting CFO responds directly to which painting is closest to the user, and scores all items of content that are relevant to that painting. The Interest CFO also responds indirectly to location. A Link is defined between Painting and Interest which specifies a number of possible values for Interest, depending on the time that Painting has held its current value. The longer that Painting holds its value, the higher the value of Interest.

If a visitor retraces their steps, the context architecture is able to determine the last known value of Interest, by consulting an internal database that stores sets of values of the CFOs. Using Painting as the search key for this database, the context system can determine what level of interest was reached last time the visitor was at this painting.

The functionality described here could be achieved using a far less involved set of rules. However, our implementation offers a high degree of flexibility and also the chance for non-programmers to easily create context aware experiences without having to worry about the specifics of the code behind the system. The gallery experience is just one example of a relatively simple application that can be deployed using our architecture. The architecture itself is designed to be flexible and extensible, to allow for much greater complexity than was used for these initial gallery trials.

We are using a bespoke ultrasound tracking system to determine the location of users as they move around the gallery space. This system has been developed at the University of Birmingham as part of another project [Cross, Wooley, Baber, & Gaffney, 2002], and has been successfully adapted to provide input to our context awareness system.

The ultrasound system comprises a set of transmitters placed at known points on the walls of the gallery, and a receiver which connects to a PocketPC device. The receiver is able to triangulate its position from the signals received from the fixed transmitters.

4 Results of user trials

From December 2004 to April 2005, we conducted user trials of our context aware system at Nottingham Castle Museum gallery. At time of writing, our results are at a preliminary stage, and we have not yet analysed data gathered from our questionnaires or audio/video recordings.

All participants were visitors to the Nottingham Castle Gallery who were approached and asked if they wished to take part in our study. All were given a brief introduction to the system and its aims. All participants (except the control condition) were asked to complete a pre- and post-task questionnaire so that we could assess what they had learned from their visit.

We gather data from several sources for our trials:

- pre- and post-task questionnaire data, to determine what visitors have learned from their visit
- video recordings: of visitors' movements in the gallery
- audio recordings: of visitors' conversations whilst using the system
- system logs: of content delivered, movement between paintings, options selected on the PDA

We used an independent measures experimental design to determine the impact of the use of our handheld guide (experimental condition) in comparison with traditional guide materials (baseline condition: a printed booklet) and no provision of guide materials at all (control condition).

Preliminary results aredrawn from informal observations taken by the experimenters during the trials.

We found that visitors using the paper guide tended to follow a more 'rigid' pattern of movement around the gallery, visiting paintings in a specific order, then stopping to consult the guide book. In contrast, visitors with the PDA were more likely to move around the gallery according to what interested them, after scanning the room for paintings that caught their eye. It seemed that because the handheld guide had no inherent structure, this structure was not imposed on the visitors' behaviour.

A number of specific problems were observed when people were using the system. People quickly developed high expectations of the system based on previous experience, often remarking on paintings that did not offer the same depth of content as the others. Content availability was apparent from the screen display, but this seemed non-intuitive for many users. Even the basic system was perceived as overly complex by many users, emphasising the need for content delivery systems such as this to remain as simple as possible. Despite perceiving the system as complex, most users seemed to find the system useful once they had discovered what it could provide

However, few users made use of the content navigation options on the device, and were content to simply have content delivered in the order the system dictated

5 Conclusions and Next Steps

The system described here has been deployed in Nottingham Castle Museum gallery and is currently undergoing user trials. Preliminary testing of our prototypes has indicated that are important research issues surrounding the use of context sensitive architectures both to drive applications and to provide alternative means of content navigation for users. The main challenges in this area are those of determining appropriate ways to represent these new metaphors for navigation to users, and creating usable interfaces within the constraints imposed by the design of mobile devices.

In particular, it seems that context-aware applications must be simultaneously opague – in the sense that the user can use the system without being concerned with the details of how it is performing its task – and optionally transparent, so that the user can inspect the state of the system, correct mistakes, and use the contextual information for their own purposes such as content navigation.

Acknowledgements

This research is supported by the EU 5th Framework Programme, and has been conducted as part of the IST Project MOBIlearn. We would like to gratefully acknowledge the contributions of non-author project partners to the work reported in this paper, and we would like to extend our appreciation to the staff at Nottingham Castle Museum who have supported this work.

References

Abowd, G. D., Atkeson, C. G., Hong, J., Long, S., Kooper, R., & Pinkerton, M. (1997). Cyberguide: a mobile context-aware tour guide. *Wireless Networks, 3*(5), 421-433.

Baber, C., Bristow, H., Cheng, S., Hedley, A., Kuriyama, Y., Lien, M., Pollard, J., & Sorrell, P. (2001). *Augmenting Museums and Art Galleries.* Paper presented at the INTERACT '01.

Beale, R., & Lonsdale, P. (2004). *Mobile Context Aware Systems: the intelligence to support tasks and effectively utilise resources.* Paper presented at the MobileHCI 2004.

Chen, G., & Kotz, D. (2000). *A survey of context-aware mobile computing research* (TR2000-381). Hanover, NH: Dartmouth College.

Cross, J., Wooley, S., Baber, C., & Gaffney, V. (2002). *Wearable computing for field archeology.* Paper presented at the International Symposium on Wearable Computing Applications, Zurich.

Dourish, P. (2004). What we talk about when we talk about context. *Personal and Ubiquitous Computing, 8*(1), 19-30.

Hsi, S. (2003). A study of user experiences mediated by nomadic web content in a museum. *Journal of Computer Assisted Learning, 19*, 308-319.

Naismith, L., & Smith, P. (2004). *Context-Sensitive Information Delivery to Visitors in a Botanic Garden.* Paper presented at the ED-MEDIA: World Conference on Educational Multimedia, Hypermedia and Telecommunications, Lugano, Switzerland.

Oppermann, R., & Specht, M. (1999, Aug). *Adaptive mobile museum guide for information and learning on demand.* Paper presented at the International conference on human-computer interaction, Munich, Germany.

Proctor, N., & Burton, J. (2003, 19-20 May 2003). *Tate Modern Multimedia Tour Pilots 2002 - 2003.* Paper presented at the Mlearn 2003 Learning with Mobile Devices, London.