Pervasive Adaptation for Enhancing Quality of Experience

Arvind Bhusate and **Jeremy Pitt**[1]

Abstract. A visit to a public collection is potentially one of the most effective and entertaining ways of knowledge acquisition open to the general public. Such collections are nevertheless under increasing pressure, for various commercial and technical reasons, to improve the qualaity of experience (QoE) afforded to users. In this work, we propose to (try to) enhance QoE in a visit to a museum through a combination of a sensor-saturated environment with intelligent decision-making and policy-based computing, to deliver personalised, context-sensitive services to the visitor. The fusion of sensor data and user actions to adapt the behaviour of, for example, displays and exhibits is addressing a specific challenge of pervasive adaptation, i.e. how to coordinate the complex interaction of users, devices, and sensors to deliver an improved service or experience.

1 INTRODUCTION

A visit to a public collection, such as a museum, is potentially one of the most effective and entertaining ways of knowledge acquisition open to the general public. Such collections are nevertheless coming under increased pressure to improve the quality of experience afforded to visitors. The reasons for this are due both to competition from alternative forms of entertainment, and also from the expectation levels raised by those forms. We believe though that the appropriate response is no to attempt to make the visit a comparable 'experience', but instead to capitalise on potential opportunities in exhibition delivery: with respect to the visitor-museum relationship, the visitors-museum relationship and finally the visitor-exhibit relationship [3].

With respect to this latter relationship, in any visit there will be those exhibits in which the visitor is more interested, but can't interact with because restrictions are in place, e.g. it is behind a glass cabinet. Even if there is any form of interaction or direct experience with the exhibit, it is currently un-recorded; yet this direct experience is potentially of most value to both the museum and visitor. The visitor's enjoyment, engagement, learning and recall are all at risk as the dialogue between the visitor and exhibit is limited to mainly viewing and reading. We need to find ways in which this limited dialogue can be expanded.

Possible ways of enriching this dialogue include tracking visitors to provide personalisation, encouraging visitors to browse and explore the exhibits in a meaningful fashion, and finding an appropriate trade-off between directed, structured or didactic routes through the exhibits, and open, unstructured paths which end up with the same pedagogic content. There is also a requirement to make subject matter tangible: in addition to digital content (which would be readily available over the internet anyway) the visitors have to be exposed to the 'real things'. We propose to 'follow' the visitors in their passage through the museum, complement digital information and/or virtual exhibits with real life replica models, and personalise the interaction, and so improve the overall quality of experience accordingly.

We propose to use a combination of pervasive computing and policy-based computing to (try to) enhance *quality of experience* (QoE) in a visit to a museum, through a combination of a sensor-saturated environment with intelligent decision-making with respect to deontic policies, in order to deliver personalised, context-sensitive services to the visitor. In this paper, we begin by offering an appropriate definition of QoE; specify the architecture of a platform to provide, measure and evaluate QoE; and describe the decision-making support which adds the 'intelligence' that in its interpretation of sensor data and user activity adapts the display and exhibit behaviour in order to enhance QoE (cf. [9]).

We then describe the iCars Exhibition, a combination of a sensor-saturated environment with intelligent decision-making and policy-based computing, to deliver personalised, context-sensitive services to the visitor. The fusion of sensor data and user actions to adapt the behaviour of, for example, displays and exhibits is addressing a specific challenge of pervasive adaptation, i.e. how to coordinate the complex interaction of users, devices, sensors, and intelligent systems to deliver an improved service or experience.

2 QUALITY OF EXPERIENCE

2.1 Background

Quality of Experience (QoE), sometimes also known as "Quality of User Experience", is a subjective measure of a person's experiences with a place, object etc. More precise definitions include:

- the degree to which a system meets the target users tacit and explicit expectations for experience [2];
- the measured level of quality of a particular user experience when compared to a specific target, using a specified metric and method or tool [8].

An experience could be of watching a video over the internet, dining in a restaurant and so on. User experience is context-dependent and tends to include wider human experience dimensions such as pleasure, fun, and other emotions [2]. In this work, a user's experience is expressed as emotions, attitudes, thoughts and perceptions felt by users across the usage life cycle and which affects the users behaviour.

Many other studies and research have been carried out on trying to understand, define and measure such a subjective area of study. Beauregards et al [2] describes Quality of Experience by breaking

[1] Department of Electrical & Electronic Engineering, Imperial College London, UK, SW7 2BT; email: {a.bhusate,j.pitt}@imperial.ac.uk

it down into two distinct parts, the first being the User Experience and the second being User Experience Quality. User Experience is described in the form of a loop, which is made up of three major components: perceptions, emotions, thoughts & attitudes and finally behaviour.

Alternatively, Alben [1] defines several criteria, which include: understanding of users, learnable, usable, needed, mutable, effective design process, appropriate, aesthetic, and manageable. Corrie et al [5] discusses four domains, which include: task domain, needs domain, services domain and technology domain. Finally, Forlizzi and Battarbee [7] uses five aspects of experiencing a product, which are: physical, sensual, cognitive, emotional and aesthetic.

2.2 QoE for a Public Collection

All of these studies are informative, but demonstrate that a definition of QoE is highly context-dependent. To derive a workable definition for our puposes, we need to

- Define the experience: This involves describing a typical yet detailed scenario (situated in an exhibition space), which aids to set the physical environment, describes the sorts of exhibits we envisage the visitor will encounter and describes the type of visitors this experience is aimed towards.
- Understand the experience: By identifying the roles, goals, relationships and interactions of the, and between the, visitors, objects and environment we can identify the factors which can affect their experience.
- Design and realise the experience: Using the above understanding we can create a requirements specification of the components which need to be developed to realise this experience.
- Measure the experience: Finally using the factors we can identify the measures/criteria and create a hypothesis to evaluate against.

In the following sections, we briefly consider each of these steps.

2.2.1 Define the Experience

We envision a system in which there is a sensor-saturated environment, where the environment is equipped with cameras, the visitor is equipped with an indentifying device, and the exhibits equipped with a range of sensors to determine users' actions, e.g. those which change the state, orientation or movement of the exhibit. This envisionment is illustrated in Figure 1, and foreshadows the iCars Exhibition presented in Section 5.

2.2.2 Understand the Experience

Assuming this 'experience envisionment', which is quite abstract, we derive from it more concrete components: visitor, exhibition space and exhibits. For each we identify their roles, goals, relationships and interactions amongst and between one another respectively. Finally the identify the measurable factors constituting QoE.

Visitor. This exhibition space is designed and created for children (both male and female) between the ages of 11-13. It addresses issues which are dealt with in school curriculums today so that the children can relate and connect with the information. The children have the opportunity to browse the exhibition and explore the exhibits. The children are empowered by a museum (in this case the exhibition space we have defined) which defines what they can and cannot do (using powers and permissions, the basic building blocks of deontic policies) with regards to interactions with the different exhibit types.

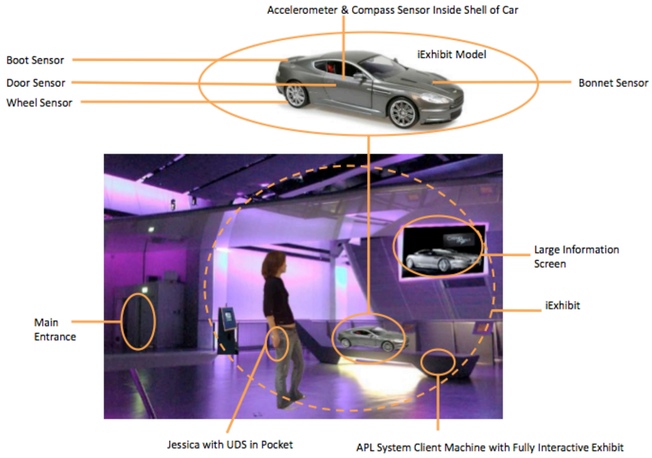

Figure 1. Interactive Exhibit

Exercise of these powers and permissions is dependant on their usage and behaviour. The children also have an opportunity to interact with one another to obtain a shared experience, however we do not model or define any governing mechanisms for this.

Exhibition Space. The exhibition space is browsed by the visitors. The aim of the exhibition is to educate children in an entertaining manner about the history and culture of cars and how they connect to important issues of security, safety & security and environment. The exhibition space houses many types of exhibits and has an overall quality associated with the many experiences it provides. A museum housing an exhibition space has the responsibility and obligation to manage and organise exhibits and visitors.

Exhibits. The exhibits are explored by the visitors. Amongst the multiple types of exhibits there are actually three main types we are interested in evaluating the QoE. Each exhibit may be traditional (non-interactive), partially interactive and fully interactive. Thus each exhibit provides an experience and has an associated quality. Exhibits have a responsibility to serve content to a visitor. However, certain exhibits can be empowered which gives them options of how it can serve its content based on the visitors empowerment levels, actions, usage and behaviour. For example an exhibit may decide that a users behaviour has been excellent and therefore should be rewarded. This would actually mean that the exhibit itself decides to increase the permission levels for the user as it has gained trust. Since this is based on trust it can potentially put the security of the exhibit and other similar exhibits at risk.

Factors. Using this scenario together with the identified visitors and objects, the relationships and interactions amongst them we define the factors which could affect the several smaller experiences present and the overall experience quality for a visitor. This is an essential step in realising and evaluating this experience. In this context, we determine that the quality of experience is made up of three main groups, which are:

- human perception factors: ambience, aesthetics, comprehension, content, accessibility, tangibility, interactivity, usability, and collaborativity;
- technology factors: functionality, personalisation, security, trust & privacy, response time, synchronisation, decision, intelligence, latency, and errors;
- human psychological factors: comprehension, emotions, stress, anxiety, patience, fatigue, behaviour, and expectation & demand.

2.2.3 *Design and Realise the Experience*

The substance of this section is described in the following sections, on the system architecture (Section 3), decision support (Section 4), and their incorporation in the interactive iCars Exhibition (Section 5).

2.2.4 *Measure the Experience*

The measurement of experience, according to factors identified in the three main groups above, is an outcome of the evaluation trials and experiments, and is addressed in Section 5 below.

3 SYSTEM ARCHITECTURE

This section presents the overall system architecture, called the APL platform. It is generic in design, but is targeted at a specific exhibition, i.e. the iCars Exhibition. This comprises a server and four constituent exhibits: three cars, one non-interactive, one partially interactive, and one fully-interactive; and an interactive quiz exhbit. However, the APL platform is generic and requires the designer to instantiate the APL platform for a specific planned exhibition.

3.1 APL Platform: Logical View

The APL Platform's purpose is two-fold. Firstly, it presents a means to provide an experience and capture this experience. By offering a richer experience it aims to of enhance the quality of that experience in a museum, by providing interactivity through the unobtrusively embedded sensors and a mechanism to dynamically deliver information in various forms governed by user/system policies and intelligent decision making.

Secondly, it is used to help evaluate the users experience by collecting, processing and storing user interaction data later to be retrieved for analysing their experience. The APL Platform keeps track of the users exploration in terms of the exhibits they have viewed, the duration they have spent at each, their focus at the time of exploration, and the interactions which took place between the exhibit and them. It also acts as a guide to aid the user in their learning. This functionality allows the user to focus on just exploring and learning and not worry about what they have seen, what they still need to see etc. This opens avenues for building a longer lasting relationship between the museum and user as they can continue their visit from where they left off.

Therefore, at a high level of functional abstraction, the APL Platform provides a user with a variety of personalised services and interaction mechanisms in a ubiquitous computing environment. It provides the administrator a means to monitor and evaluate a users experience in an unobtrusive manner.

The APL Platform is (logically) made up of software components, collectively known as the APL System, and hardware components, which in the iCars Exhibition include the APL Exhibit Extension, ID Devices, video cameras and quiz pads. The APL System software consists of two parts: the server-side component, namely the APL System Server and the client-side component, namely the APL System Client. The client side entity directly provides the personalised services and content using its decision-maker in conjunction with fuzzy inference systems plus defined policies (see the next section). However the server side component is working together, along with other user-defined policies to cater for this personalisation of services and content. The APL System Client communicates with the APL System Server and the Exhibit Extension (if any) using Bluetooth.

3.2 APL Platform: Physical View

The *physical* architecture of the APL platform, i.e. in terms of machines, for the iCars Exhibition, is illustrated in Figure 2. Recall there were four constituent exhibits, therefore one machine for each, each running a separate instance of the APL System Client, and each with its own local policies and content database, display screen to present information about the exhibit which it is associated with in the form of text, images, videos and CG models, dependant upon the type of exhibit, video camera embedded in the display screen (to be unobtrusive) which captures the user's focus and attention (its operation is automated and is triggered by the users presence). In the case of Exhibits 3 and 4, which are interactive, there is an Exhibit Extension which includes all the sensors and hardware required to implement an interactive exhibit. The visitor carries an ID device (IDD) which is used to make the various components of the platform aware of the presence of whoever is logged into the system with it. It is carried by a registered user to identify their presence. The IDD sends data via radio to the APLSystemClient.

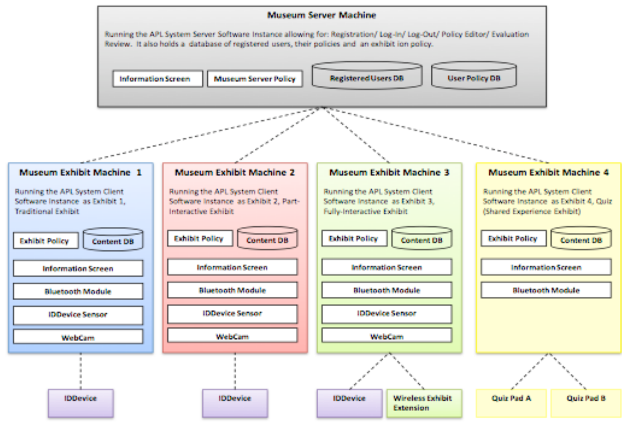

Figure 2. iCars Exhibition Physical Architecture

In more detail, the function of each APL Platform (software and hardware component) is as follows.

The APL System Server allows for user registration, user log-in/log-out, and user/system policy creation and editing with a built in policy configuration toolkit. Registration allows a user to register to the APL Platform, where an account (folder and UserID) is created for them. This account initially stores a newly created policy and will hold future SessionInteractionData and VideoData files. Along side this, their name, DoB and UserID are also stored in a Registered Users Database. Log In simply associates an IDDevice to a user and broadcasts their policy over the Bluetooth Communication Link to all available exhibits allowing them to start exploring. Log out simply disassociates a user from an IDDevice after they have finished exploring and makes the IDDevice available to another user.

The Policy Configuration Toolkit allows for the creation and modification of user/system (exhibit) policies catering for the personalisation and added value of the service received by the user. The APLSS also manages and monitors Exhibits (APLSystemClient instances) and their policies.

An APL System Client instance is able to detect the presence of users that are registered and logged in via their IDDevice. The APLSC is receives user policies from the APLSS and then loads them ready to be used when the corresponding user is detected. It creates a session for this detected user and then captures the users interactions (particular to the fully-interactive exhibit) and video data. It also provides information using a variety and mixture of multimedia content including a virtual CG model dependant on the exhibit type. The automation and control of these processes is carried using an intelligent decision maker which uses both the user and exhibit policies.

APL System Client Software retrieves and loads appropriate user policies dependant on the user(s) detected. It also collects data from the user exhibit interactions, processes it and encapsulates it into the users session file. After the user leaves the exhibit, their policy and session file is sent to the APL System Server Software. It also makes decisions about what content to serve and when, based upon the users actions, overall behaviour and several policies including user and exhibit policies. It provides information content about the exhibit to the user via its associated information screen.

4 DECISION SUPPORT

4.1 APL System Client 'Intelligence'

We aim to enhance QoE through some 'intelligence' in the fully-interactive exhibits. This is derived from fusing the behaviours of users derived from sensors with other data, using a fuzzy inference system (FIS). The output of the FIS, in conjunction with user policies, is interpreted by the Decision Maker, for example to change the view on the associated display. This is illustrated in Figure 3.

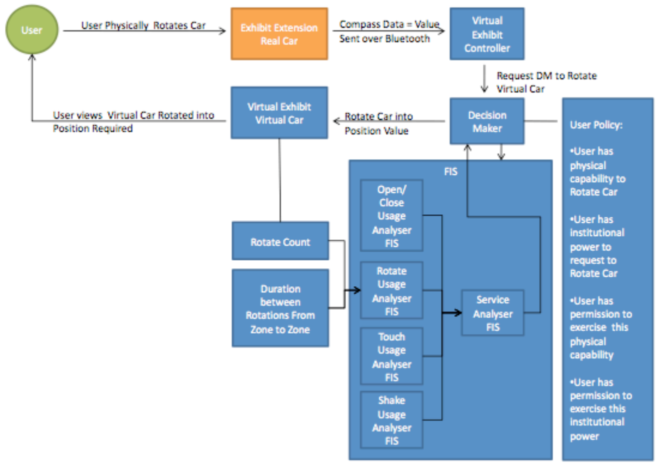

Figure 3. Use Case for User 'Rotate Car' Action

In the rest of this section, we review the three main components of the decision support system: the APL policies, the fuzzy inference system, and the decision maker.

4.2 APL Policies

An APL Policy can be written for several actors of the APL Platform. These policies encapsulate setup information for each entity. They also indentify and define the responsibilities in terms of obligations and sanctions for the system components; and actions in terms of physical capability, institutional power and permissions for the subjects (users). In particular:

- a policy represents the museum server in terms of its setup and main responsibilities, e.g. delivering system and subject policies as and when requested, keeping track of users movements etc.
- a policy represents the exhibit in terms of how it must be set up, its responsibilities which importantly govern what reactions to carry out on certain actions from the user.
- a policy represents the user in terms of the actions they can/cannot perform whilst browsing or exploring an exhibit. A user can only have one policy file associated to them.

There are a number of alternative policy languages available (e.g. [6, 10, 11]), but these were aimed at different domains (management, security, web services), none quite captured this full range of representation, and did not distinguish between being physically capable, being conventionally empowered to bring about (his/her actions had conventional significance in the context of the museum), and whether or not the user was permitted to perform either physical or conventional acts.

For this reason, we designed our own policy language, based on a museum/exhibit ontology of actions and objects, policy information model, dictionary and grammar, collectively defining the APL Policy Language.

This language is used to write APL Policies. describe and in turn govern the operational (functional) behaviour of an APL System Client (Exhibit). This will allow for control of the content that will be presented to a user, helping to protect the exhibit from any potential harm etc.

There are currently two types of APL policies: subject policies and system policies. The subject policies contain details about the users (visitors or administrators), and also defines their access control for applying actions on exhibits. The details of the visitor include their personal information and the actions they are permitted to apply using the concepts of physical capability, institutional powers and permissions to exercise the former two. The system policies contain details about the various system components, and defines the systems control for automation of tasks such as delivery of content, denial or service, warning or aiding between system user and system system type interactions. The system uses obligations and sanction policies to control the delivery of content to the user, the general running of the system and finally what sanctions to apply when the user breaks a policy. The elements of either policy type can be specified in any order. The elements of either policy types can contain many other policy types. A system has many default system policies. Some system policies can be overruled if stated in the system policy by an administrator policy for example. Access control is a good example of this.

We use in our policies the distinction, which is standard in the study of legal, social and organizational theories, between three types of 'can'. This is the physical capability to manipulate objects (e.g. manually rotating an object), the institutional power to establish conventional facts (e.g. a command to rotate an object), and the permission to perform either type of action.

- *Physical Capability (Can_1) Policy*. We assume the user always has physical capability. However, they might not have permission to pick up and perform a physical action (rotate, touch, open/close etc) on the physical model of the exhibit if they do, it could sound an alarm alerting a nearby member of staff to attend the scene. Or the virtual exhibit could simply not respond leading the user to boredom.
- *Institutional Power (Can_2) Policy*. Power to request to rotate the CG model of the car (create a conventional fact), whether these requests are acknowledged and served by the system depends on if the user has the permission to exercise this power, if so the model will move appropriately. If this power is not granted the requests

cant be made, however will be logged as part of the data collection. The CG model will be frozen.

- *Permission (Can_3) Policy*. The Museum gives the user permission to exercise his/ her physical capability and institutional power on the physical exhibit in order to control the virtual exhibit i.e. open/close, rotate, touch and shake.

Permission can apply to both the physical capability and the institutional power. You can have permission to make a request, and you can have permission to pick up an object, or not, in both cases. The difference is, we can intercept in the system itself, the institutional power and permission, and grant/withdraw these accordingly. If the user has permission to make a request, a rotate request for example, the request that is made will be honoured and the relevant information will be provided (rotation of the CG model will occur, textual information will be displayed, video will start etc. at appropriate positions (zones) the model is rotated into). If the user does not have the permission to be able to perform *can_2*, the request won't be honoured and the CG model will appear frozen.

We can *not* affect the physical capability of the users (short of "getting medieval" on transgressors); we can however monitor and control the permission to physically interact with exhibits, but can only do anything about it if we alert a Docent (i.e. back in the "real world", outside the system itself).

We will grant institutional power to specific users one at a time, so their movements do actually count as requests to the controller to perform actions of the appropriate type. This will also guide the user through the experience that was intended (so as for the user to systematically learn and query the exhibit) rather than just letting the user randomly explore it.

4.3 Fuzzy Inference System(s)

Fuzzy Logic is a formalism which facilitates reasoning about imprecise facts, uncertainties, and value judgments. Fuzzy Logic is the basis of fuzzy inference systems, although there are different types of fuzzy systems as there are various different ways in which outputs can be determined.

In general, to build a fuzzy system, an engineer might start with a set of application- dependent fuzzy rules as specified by a domain expert. Fuzzy rules are expressed in the form "if . . . then . . ." that convert inputs to outputs, where both inputs and outputs are members of fuzzy sets (a fuzzy set is a set in which objects are members to some degree).

Given a set of such rules, it may be that a particular range of inputs fire (activate) any given subset of those rules. The rules which are fired then contribute proportionally to the fuzzy output: this is calculated by applying the implication method of fuzzy logic to the activated rules and aggregating all the results. The process of defuzzification converts the aggregated output into a 'crisp' value (the usual method is a centroid calculation, i.e. finding the centre of an area under a curve).

This entire process, called fuzzy inference, thus converts quantitative inputs into a precise output using qualitative statements.

Several FISs are used to provide the user with a service (present information to the user), help and finally protect the exhibit if negative intentions are sensed by inferring a users usage type and behaviour. The FIS is used to determine if: the user needs help, is mistreating the exhibit, or should be granted permission to be able to apply more advanced actions.

If the user is determined to be mistreating the exhibit or not adhering to boundaries specified by instructions or warnings etc, the FIS will sanction the user. The severity levels of the sanctions are in the following order (1 being lowest level of severity):

1. Warning of removal of permission
2. Removal of permission
3. Warning of removal of power
4. Removal of power

For example, a user has the physical power to touch the steering wheel, they have the institutional power to touch the steering wheel and request information about it. However, they do not have the permission to touch the steering wheel when a door is not open. In this case the user will just be helped with a message such as: this is not allowed, please open a door first. If this was the users first minor offence and is not doing any harm, then the FIS doesn't log this as a problem. If the user does this again, the FIS will see this as a warning should be provided as still no harm is done. But if the user decides to do this again the FIS will apply a heavier sanction and remove the permission of being able to perform Touch actions at all.

4.4 Decision Maker

The DM uses several APL policies and FISs to make decisions about how the system (exhibit) should be controlled to cater for the user in the best possible way as to keep the users QoE high and to keep the systems (exhibits) safe from misuse. The policies include MuseumSystem, exhibit, subject, exhibit extension and VirtualModel policies. The FISs include TouchAction, ShakeAction, RotationAction, Open/Close usage FISs. The decision maker of course gets raw data from a variety of sensors such as user presence, compass, accelerometer, touch, open/close sensor readings from the exhibit extension, time etc. to infer or process data to something meaningful which can then be used to make decisions.

The combination of FISs work together to infer behaviours, such as if the user is 'normal' or whether they are potentially 'dangerous' and misusing the exhibit. This inferred behaviour may actually not mirror the reality, i.e. this misuse may not be intentional and therefore the system should not penalize the user on just some of the interactions and therefore must not appear to be a rigid system. The FISs use multiple rules to work out from many actions if the user is in need of help or should be heavily sanctioned for dangerous/mischievous behaviour. Therefore the feedback/response from the system varies from help and mild warnings, to strict warnings and even sanctions (denial of further interactions and/or even complete removal from using the exhibit). The administrator is able to intervene at anytime with the decisions of the system if something is clearly going wrong.

5 THE iCARS EXHIBITION

To evaluate the enhancement of QoE, we have built a usability room that is a microcosm of an exhibition space in the Science Museum, London. This is the iCars Exhibition with the four exhibits, History, Security & Safety, Environment and Quiz, arranged as shown in Figure 4. 'Visitors' start in the south-west corner and view consecutively the History (non-interactive) exhibit, the Safety (partially interactive) exhibit, and the Environment (fully interactive) exhibit, before addressing the Quiz (fully interactive) exhibit.

The Environment Exhibit is an APL Exhibit Extension of an instance of the APL System Client. The exhibit itself is made up of two parts: real and virtual. The real part of the exhibit is an Aston Martin DBS scale model, which is fitted with a variety of sensors:

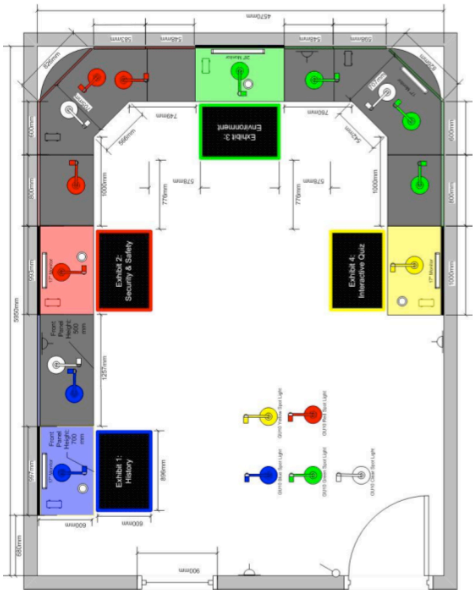

Figure 4. Floorplan of the iCars Exhibition

- Touch sensor: identifies which non-moving parts of the exhibit are being explored;
- Door/bonnet/boot sensors: identify which moving parts of the exhibit are being explored;
- Compass sensor: identifies the position of the car in the Z plane;
- 3-axis accelerometer sensor: identified movement of the car as a whole in the X, Y and Z planes.

These sensors record touch, rotations, open/close and shake actions. Detected actions can be fused with data in order to control the content delivery to the user using the information plaque, this being a CG model of the car (i.e. the virtual part). The car itself, and corresponding displays, are illustrated in Figure 5.

Figure 5. Exhibit 3: Model Car and Display

We are now in a position to commence trials and experiments to determine any enhancements of QoE enabled by the iCars Exhibition. These trials are scheduled to take place in March 2009.

6 SUMMARY AND CONCLUSIONS

In summary, we have observed that one of the most pressing problem facing modern museums is to engage their visitors in a way which is more interactive than traditional exhibits and information displays; yet is more engaging that just viewing the a computer-generated version of the exhibit on a web page, say; and yet remains informative, instructive, and memorable. We defined a notion of Quality of Experience (QoE) in this context, and developed a system designed to enhance QoE in a visit to a public collection.

This system operates in a sensor-saturated environment, with sensors in the exhibition space and in the exhibits themselves. In addition to this the client-server APL platform provides a generic inrastructure to define exhibits. Then, by using fuzzy inference and user-defined policies, we could use an 'intelligent' decision maker to customise displays and exhibit behaviours. It was this personalisation that we though would be the keystone to enhancing QoE. We have defined and built an experience, iCars Exhibition, a microcosm of a real museum exhibition space, which at the same time allows us to evaluate any enhancement.

It is too early to draw any conclusions, as the research programme is entering its evaluation phase. The system will be evaluated as it currently stands, but one direction of further work is to enhance the decision-making with a 'learning' component. The idea here is to match a current trajectory of behaviour against a history of previous interactions to create a sense of 'expectation', which can be used to provide some form of pro-active behaviour [4]. We would claim that this fusion of sensor data and user actions to adapt the behaviour of, for example, displays and exhibits is addressing a specific challenge of pervasive adaptation, i.e. how to coordinate the complex interaction of users, devices, sensors, and intelligent systems to deliver an improved service or experience.

ACKNOWLEDGEMENTS

We would like to thank the EU DANAE Project, the EU PANORAMA Coordination Action, and an Imperial College Studentship for various aspects of support for this research.

REFERENCES

[1] L. Alben, 'Quality of experience: defining the criteria for effective interaction design', *ACM Interactions*, **3**, 11–15, (1996).

[2] R. Beauregard, A. Younkin, P. Corriveau, R. Doherty, and E. Salskov, 'Assessing the quality of user experience', *Intel Technology Journal*, (2007).

[3] A. Bhusate, L. Kamara, and J. Pitt, 'Enhancing the quality of experience in cultural heritage settings', in *Proc. 1st European Workshop Intelligent Technologies for Cultural Heritage Exploitation, 17th European Conf. Artificial Intelligence*, pp. 1–13, (2006).

[4] C. Castelfranchi, R. Falcone, and M. Piunti, 'Agents with anticipatory behaviors: To be cautious in a risky environment', in *ECAI 2006, 17th European Conference on Artificial Intelligence*, eds., G. Brewka, S. Coradeschi, A. Perini, and P. Traverso, pp. 693–694. IOS Press, (2006).

[5] B. Corrie, H. Wong, T. Zimmerman, S. Marsh, A. Patrick, J. Singer, B. Emond, and S. Noël, 'Towards quality of experience in advanced collaborative environments', in *Third Annual Workshop on Advanced Collaborative Environments*, (2003).

[6] N. Damianou, N. Dulay, E. Lupu, and M. Sloman, 'The ponder policy specification language', in *POLICY*, pp. 18–38, (2001).

[7] J. Forlizzi and K. Battarbee, 'Understanding experience in interactive systems', in *DIS '04: Proceedings of the 5th conference on Designing Interactive Systems*, pp. 261–268. ACM, (2004).

[8] T. Hall and L. Bannon, 'Designing ubiquitous computing to enhance children's interaction in museums', in *Proceedings of the 2005 conference on Interaction design and children*, pp. 62–69. ACM, (2005).

[9] A. Isard, J. Oberlander, I. Androutsopoulos, and C. Matheson, 'Speaking the users' languages', *IEEE Intelligent Systems Magazine*, **18**(1), 40–45, (2003).

[10] L. Kagal, T. Finin, and A. Joshi, 'A policy language for pervasive systems', in *Fourth IEEE Int. Workshop on Policies for Distributed Systems and Networks*, (2003).

[11] A. Uszok, J. Bradshaw, M. Johnson, R. Jeffers, A. Tate, J. Dalton, and S. Aitken, 'Kaos policy management for semantic web services', *IEEE Intelligent Systems*, **19**, 32–41, (2004).