

The HyperMuseum

Stephen Jones
387 Riley St Surry Hills
NSW 2010 Australia
sjones@culture.com.au

This paper explores the greater development and function of museums. It looks at ways to interrogate a museum's collection and knowledge database according to the needs of the user, dynamically customising various aspects of the museum's role in, for example, the systems for presentation of objects and interpretation to the user. Means for making the museum aware of visitor enquiries and visitor tracking are proposed, along with the potential for rendering the museum predictive in its interaction with the visitor. Finally some indications of the means by which the museum as an environment may be made intelligent, based on recent work in robotics are suggested.

Introduction

What will the museum of the future look like and how will it behave? In this era of the Internet and hypertext Vannevar Bush's projections of hypermedia [1] have been realised, albeit in technologies that he might be hard pushed to recognise. His ideas are the roots of our cyberspatial world, yet we are remote, isolated from the object or discussion, having only the virtual object of somebody's well-intentioned editing for our experience. So can we re-ground this in the *actual* while retaining the correlating usefulness of hypermedia? Can we take the physical museum and turn it into the HyperMuseum? Not the virtual shopping mall or virtual museum that the net offers but the solidly manifest, physical museum that you can walk into. What potential for interaction is there in the actual museum and how can it be enhanced?

This paper explores the possibility of producing a museum which is a physical hypertext, in which the potential for interaction and collaboration is readily enhanced by the actual museum as system in which it becomes a behaving, adaptive, predictive, learning and collaborative entity, aware of its staff and its visitors within itself. Although I have been told by several curators that I am not at all the typical museum goer perhaps the ideas that I present here can provoke thoughts into not simply how to add value to that "typical" museum experience, but into how the museum visitor might be lead into taking a more active interest in museums and what they have to offer.

Primarily a museum is a building which houses a collection of objects, the documentation and display of which is done by a group of curators who gather historical documents and records, and display, with interpretation, the fruits of this collection. Essentially the museum is a database. The data within it consists in the objects, records

associated with the objects, records of the growth and development of an idea or a technology or an organisation, books, papers, journals which detail the operations within some technology or some historical grouping, and the knowledge and skills of the people who look after the objects and the records, conservators, registrars, curators, designers, a/v producers and so on. All of these people have knowledge of and insights into the objects, their relations to other objects, the history of the field to which the objects belong the cultures in which particular objects were produced and so on. All of this information is important for the person who goes to the museum to find out about something, or to look at a collection of computers or fabrics or sculptures, or the family who go for an interesting outing with the kids.

Extending the Database

Presently, the museum's database is accessible only to the people who work within the museum and in many ways is really operated by the registrars who record what the object is, when it was acquired, from whom and, hopefully, when it was made and why. But the museum also includes the knowledge and skills of its staff. What, of the curator's knowledge, is in the database, stored in the memory of the computer? What of the conservator's skills in refurbishing the object, what of the original user's skills in the object's use, what of the previous owner's knowledge of the history of the particular item? Why is this item important whereas another, similar item, is not so important? Why is the item culturally important at all? All this sort of knowledge is what gets poured into the display process, as installation and design, as labelling and a/v production, but then when the exhibition is finished then what of it? Into the catalogue one hopes, but it should also go into the database so that others can access it later, not as some edited text but in the full richness of the experience that went into the object in the first place. The experts need to teach the system, not so that they can be replaced but because they can't be everywhere at once and eventually they will move on anyway.

But while the database carries information, it is not experienced knowledge. What the visitor needs is to be able to develop this information into experienced knowledge: the sort that the original user of the object might have had. The knowledge that is often described as being in one's bones, muscle memory, the kind of knowledge that allows one to get back onto a bicycle some twenty years since the last time and with very little effort

be riding again as though there had never been a long break at all. Is it possible that the museum can be brought from being merely a data base to being something which embodies the direct experiential knowledge of its collected objects and might be able to impart that knowledge to its new visitors?

So we have a database. Can a database be made to actually store all this? Even if we do store it all how can it be made accessible? How will the visitor access the information they want? How do we deal with the fact that no two individuals are going to enter the museum with identical desired trajectories of exploration? And most importantly how can the visitor walk away again with actual knowledge, experience of the object as history and as production? This paper offers some suggestions about making knowledge and objects accessible.

Housing the objects

One of the primary problems is housing the objects and the decisions about which objects to show. I want to suggest that the current trend of showing the minimum number of objects in some designed space where the design overrides the collection is not working. It puts me off and I suspect that it will be putting a lot of others off, too. Why do I go to a museum? I want to see the objects and understand how they work and were made and what their relations are with other similar and earlier or later objects. I want to hear the stories of the people who made or used them and I want to see and hear what they did with the object, where it fits in the culture. I might even want to know how to produce a similar object or, where the object is a tool of some sort, how to use it to do something that I might be working on. I want access to the curator's knowledge, to the original users knowledge, to the technician's knowledge of how to use or how to make it. How can I get that?

Here I suggest a layout or design change that would significantly ease up the access to all these things that I look for. Remember the Habitat store? The thing they did that radically altered shopping, display, warehousing and so much of the hidden layer in the distribution of goods for the household was that they made the shop and the warehouse the same place. The goods were all laid out in what in museum terms could be called a "mass planting". The stuff was all there in front of you. If the store had sixteen copies of this chair and 200 copies of that cutlery set then that's what you saw in front of you on the shelving and you could select whichever you wanted. Warehousing and customer choosing were concatenated into the "store" (exploiting both its meanings most effectively). So I am proposing that the museum warehouse be turned into the museum. The order of display can be according to whatever might be most appropriate, often year of manufacture within area of function perhaps within region of manufacture. For example: 1822; Babbage's first model of the Difference Engine; within calculating machines; in the computing collection or, perhaps, in the British region. (That object is currently in the Babbage collection at the Science Museum in London.)

Displaying the Museum's Skills

If the object is not in perfect condition, so what? First it might not be because it was acquired from somebody's back shed wherein it has lain for the last 50 years. It's corroded and dirty and incomplete. Supposing that it is important enough to keep (always a difficult question) then into the display spaces. The first people who work with it will be the registrars and then the curators whom, meanwhile, might have been gathering the story of its origins from the people who donated it. Then come the conservators and the technicians to make it work. All this can be done within the display area and their activities, for many a museum visitor, would in itself be a fascinating process. How do we maintain these things? How do we use these things? These questions often come to mind as one wanders through a museum.

And the skills of repairing and reconditioning are in themselves valuable skills that should not be lost. In the current era of virtualisation we are in danger of losing many of the skills that got us to this stage. What happens if we let the last person who knew how to make that kind of textile, die without passing on their skills? The textile can never be made again, and worse existing examples cannot even be properly repaired. Suddenly we are left with the game of preserving by prevention. Do not touch! The object now becomes sacred, we cannot display it because it might break down, gather dust. We lose touch with our past, with our culture, and thence with our future because we lose track of the continuity.

One example of how to avoid this problem is the Exploratorium in San Francisco, where there are many objects which can stimulate or enlighten displayed readily all over the space and the work of repair is done in visible workshops with small teaching areas attached to the display areas. But in the Exploratorium since the emphasis is on hands on experience, many of the objects are little more than toys. Another example is in the "Making of the Modern World" exhibition at the Science Museum in London. Here is a mass planting, though still of carefully selected objects and the labelling is hardly adequate for the enquiring mind, yet there is little opportunity for the visitor to make further enquiries, and although one can freely make one's own path there is little that guides one on the evolution of a technology or on the evolution of solutions to the needs of a culture.

A New Memex, or the Hypermuseum

So back to the shop. Remember the Habitat store? It has now become a store of ancient objects. Obviously some of them cannot be touched, they really are too delicate. So place them under glass, place them on mirrors so that they can be seen from all angles. Light them well but they need only be lit when someone comes up to them, or when the museum itself is designing your tour, but we'll come to that.

My proposal then, is that the museum should become a kind of large scale physical Memex machine [1] and that the traveller through this machine, the voyager on this

database would leave a trail which not only they can know but the museum itself can track and work with predictively and knowledgeable so that it can make suggestions for the voyager. So that it can lead them into areas they may not have thought about and can provide them with a record of their voyage so that it can be handed on to others who might also take an interest in the topic of the moment.

How can this be achieved? I have outlined the placement of objects and I have pointed out the sources of knowledge. In a sense the sources of knowledge as they appear in the database are the curators' and conservators' experiences in gathering and preserving the object, the librarians' results in pursuing and allocating books and documents, the registrars' notes in recording the wherefrom and the wherenow. But the visitor in this system, how can they be recognised? How can their task be divined by the museum? How can the museum present them with its knowledge so that it may teach?

Tracking the Visitor's Interests

My suggestions follow: the museum voyager arrives at the entrance and is provided with an ID card on which they record their identity, what they are interested in looking at, whether they want a detailed search or a simple introduction, and whatever else might be useful, such as whether this is their first visit, a follow up or a departure from some previous interest. The voyager is then given a card to carry which can be electronically recognised by the museum as they travel through it. If the visitor has a PDA (like the Palm Pilot or a Psion) they might communicate via this. Thus the museum knows where the visitor is within it and what they want to see. And knowing where the visitor is it can lead them to the right gallery and show them what objects they want to see.

Let's say that one has come to the museum in order to study a certain kind of fine Chinese ceramic. As one enters the building and has filled in the appropriate search form the museum indicates that one should go to a particular corridor and into a particular area where numerous ceramic objects are housed. All the ceramics in the museum's collection will be there but how do we isolate out the ones we want to know about? By lighting each object specifically (by, for example, using pin-spots, one for each object, or motorised lights which can be pointed at the current object of proposed interest) and by setting up the museum so that it knows exactly where everything is and which lights could point to what object. Then it is just a matter of switching on the appropriate light. If one wants to know what kinds of ceramics preceded this one in development then the museum can point to the appropriate objects in a sequence. It is no problem for the voyager to make requests of the museum, their ID card or their PDA will have a downloaded application presenting menus by which they can communicate with the museum and, for example, ask for a sequence or a grouping of similar pieces by different artisans of the same period, or a collection that follows its development from the original piece up to the present day. Simple wireless or IR networking will allow that communication. When the voyager has looked at all

the objects they might want to know what firing techniques were used, or where the glaze chemicals were mined or how much similar pieces fetched at auction. All of these can be easily drawn from the database and the visitor can then be directed to the appropriate gallery and shown how the kiln was built and how it works, whether this is a real kiln or an a/v of the one at the pottery.

As PDA's and next generation phones become commonplace, all the visitor will have to do is greet the museum on entry by communicating with it through their phone or PDA. On recognising the visitor it will know whether they have been here before and what they were interested in last time. When the visitor indicates their intentions this time the museum will be able to follow up on what they explored previously or if this is a new subject it will be able to enquire as to what aspect of the topic they might wish to pursue and then proceed accordingly. Of course you could simply say that you just wanted to wander through and browse freely with no particular aim in mind and the museum can light spaces as you enter them, perhaps present more detail as you linger in front of something and dim the room again as you move to the next. And for the whole family, who cares if the kids get lost, the museum will know where they are.

With a well set-up database the technical details and the collecting history, the oral history of the users, the written details of curators and scholars, the practical skills of the operator and conservator can all be linked and readily accessible. If exploring using a PDA the museum can offer the available resources, and the visitor can select what they would like presented, either direct to their PDA or onto screens by the display shelves and to nearby audio devices or directly to the visitors own headphones (with the ubiquity of walker CD players explaining the works in today's museums, this is only a simple extension). If the visitor doesn't have a PDA or similar device then, with a visitor ID card, the museum can easily accommodate by simply tracking them and presenting details and menus on display screens which can be liberally placed in the exhibition areas.

So the museum can be proactive and recommend lines to follow or it can be passive and simply wait for requests. All of the technologies I have outlined above exist today. The ID card is a readily available corporate ID card containing a Radio Frequency Identity Device (RFID) which is powered by the signal that interrogates it. PDAs, Handi-phones and walker CD players are all contemporary technology. Their communications with each other, as now, can be simply extended with the use of in-house infrared or wireless (eg. Bluetooth) communications links. Wireless is non line-of-sight and recent work in IR is also overcoming this problem. Databases built in XML can easily allow the storage and linking of all the classes of records that might be available and most museums already keep their records in some sort of database system. The display "shelving" has many available solutions and mobile lighting is used in almost every theatre or dance-club. Flat-screen displays are now common and recent developments in the focussing of sound will make the potential interference that public audio presents no longer a problem.

The Adaptive Museum

But more importantly this museum can learn and when it has experienced enough different visitors it can start to understand what people do and what they want to see, then it can start to make predictions. If it knows deeply enough what is in its database then it can actually start to tell the visitor its stories, like a grandfather telling of the family migration or your grandmother showing you how to cook *that* pudding. How can this learning behaviour develop and how can the museum be made adaptive, ready to re-jig its “conversation” with the visitor on the fly?

The database must have a capacity to not only “replay” any of its content but to tailor that content to the visitor's previous experience (as far as that is known by the system). It needs to be able to adaptively re-label objects and appropriately offer interpretive materials on the basis of the visitor's previous use of the system and the database's knowledge of what might be appropriately connected as ideas or knowledge about the object. When a system can make these kinds of judgements about the visitor then it can actually start to predict what might be the most interesting way to extend the visitor's experience. Ultimately it is this predictive capability that I am interested in.

Prediction, here, is a matter of making inferences based on existing database connections coupled with knowledge about the visitor's history within the museum and their stated intentions in coming into the museum. Prediction might be a simple mechanical statement of the available connections between knowledge items or it might be a more common-sense based inferential system in which the very database also contains a great deal of knowledge of what happens in the real world so that any predictive moves are not ludicrous (though that may be a way into creative and lateral solution generation). A database imbued with common sense is the kind of system that is being developed as CYC, originally out of the University of Texas at Austin [2]. But prediction is an overlay, an extension built from common sense and history. It comes from having some idea of what the visitor already knows. That is, it requires drawing feedback from the visitor and comparing it with the current state of the database's knowledge. Does the feedback show any gaps? What does the system know that the visitor hasn't seen yet? Does the visitor know about this? The museum enquires of the visitor and/or the visitor's feedback provides a basis for comparison and what falls out is what might well be a new avenue for exploration.

The information elements in the database, presumably stored in some sort of XML form, need to be dynamically re-formatable according to the receiving device on which that information will be displayed. So I would suggest here that it is incumbent on the receiving device, be it PDA or Handi-phone or LCDisplay, to carry out the reformatting task. Within the Handi-Phone one might want to present the information as speech and other audio, with the PDA a limited version of the website presentation would work, while within the museum's own displays text, graphics, audio and video can be comfortably handled. Obviously

appropriate combinations of these options would be best, for example the museum display provides the information while the visitor uses their PDA to give feedback to the museum, perhaps in the form of enquiries.

As the museum acquires more information about any particular visitor it might recognise that the visitor's feedback is showing connections between knowledge elements that it doesn't already have set up. Perhaps the visitor is an expert on this device or object. Perhaps they can provide information for the knowledge base? The museum should flag this event and needs to have a means of pursuing the possibility just as any curator would have. The museum needs to be adaptive and to recognise that in a sense its knowledge base is always incomplete, just as most visitors would readily acknowledge of theirs. Now the museum can learn. At this point, with current technologies this is probably about as far as the museum can go without sending in a curator to talk to the visitor but if artificial intelligence can be brought to fruit then, in the medium-term future, the museum itself could well be able to interview the visitor, logging, storing and integrating the visitor's knowledge with all the other material that is relevant within its database.

These levels of adaptability aren't strictly artificial intelligence, they are more at the level of the autonomous agent. But this kind of organisation could well provide a clear basis for the development in the system of some sort of real “intelligence”. Neither is this a disembodied intelligence. Here the system is able to know about the whole state of its body, the physical museum, as well as having any number of distributed organs in it: the galleries, the technical services and their representatives, curators, registrars, conservators, administration, etc., which give this museum internal operators carrying out the processes of self-regulation that in smaller more discrete entities would be thought of as life-defining. Yet this body does not attack its “strangers” but welcomes them in and guides and helps them, it would have quite a different kind of immune system than the bodies we generally think of as living. Perhaps these visitors are more like food, not only to be digested but to be fed as well.

Going the Next Step

Developing real artificial intelligence is not yet a simple process. It is no longer considered a matter of top-down programming or hard-wiring. Most new AI work is based on a more-or-less bottom up approach which is often seen best in the development of robotic systems, eg. Rodney Brooks' work, at MIT, in the development of behaving robots such as COG [3] or Cynthia Breazeal's Kismet [4] (also at MIT). The reason for this is that most previous work in robotics has been singularly unsuccessful, with the robot becoming quite unable to perform whenever the environment changes in a way that the robot has not been programmed to handle. Artificially intelligent robots must be able to handle unpredictable variation in the environment, the sort of thing that is way beyond anything that might be easily predicted by some sort of formal programming technique. The approach being explored

now involves building up low level, operating hardware that can handle the sensory (input) and effector (output) behaviours so that basic functions are established and brought to operational completion before the next levels are laid on top. This is largely an epigenetic process and is often given the name “subsumption architecture” following Brooks formulation of this approach for the development of COG. Epigenesis is essentially what animals and humans go through as they developmentally mature in the real world in which they live. In the growth of the infant the basic function of crawling has to be established before it can pull itself up so that it can then start to learn how to walk. Or more primitively, light sensors had to evolve before eyes could appear.

Once a basic level of function is established then the next levels can be hooked in to the system so that they receive what they need from input sensors in order to make useful adjustments to effectors in order to accommodate what the sensors are presenting. In more complex systems the intervening processes of analysis of sensory data for appropriate effector use are handled by neural networks. For a long time these were set up as pre-trained devices in simulation and were not of much value in robotic AI but they are now often used by being trained within the live real environment that the robot will be operating in. This brings a level of flexibility not previously available and may be seen in the work of Phil Husbands and Inman Harvey at Sussex University [5].

But what we are looking for in the intelligent museum is not a robot. What we want here is something more akin to a ubiquitous-computing environment in which the museum as environment becomes intelligent. As I have argued briefly elsewhere [6] there is not a great deal of difference between an environment equipped with the necessary sensors and communicative channels to allow it to be intelligent (or a robot which is similarly equipped) and a living organism except that the environment as building also houses within it an environment which is inhabited by others, people working within it and people using it as a tool. The point here is that an intelligent environment could well be developed using the same kinds of evolving neural network techniques as employed in robotics development.

And it is in this arena that the whole function of the sensory, communicative and effector-driving processes that would be needed by an intelligent museum become something that might be amenable to serious artificial intelligence development research. Because the museum is an already established database and has within it many of the training functions that neural networks need as well as housing a number of experts who can design, implement and maintain its systems it already has many of the “organs” that an artificial entity could effectively use as part of its “body”.

Establishing the right kinds of feedback structures is of course a crucial issue here. In humans we do this by being thoroughly over-endowed with wiring at birth and then by holding on to only that wiring which gets used. This is a Boltzmann network architecture in which the system starts with all connections made, and only those that are regularly reinforced by use are retained. Thus all the

connections that the system needs are available, and it can then find accommodations to any perturbations it undergoes. It is the feedback of data through the re-entrant inputs that create the attractor states in the system that the states of other parts of the network then fall into supplying a balancing process in the maintenance of homeostasis.

Ultimately the key to intelligence, as it is to living systems, is self-regulation. It is this that the feedback function provides, the means by which the system can know when something has changed within its environment. Comparison of states by the use of feedback means that it is able to know what kind of “gap” in its current state has been opened up and, given that the feedback loops have been tied into the right structures or networks, what kind of measures are necessary for the control and continued survival of the system. It is the use of feedback structures that allow all factors, human and synthetic, to compare their behaviours for completion states versus incompleteness states, and, for example, permitting the museum to predict what else the visitor might like to find out about.

Conclusion

I have, in a sense, revamped Vannevar Bush's original concept of the Memex machine as a hypertextual information collation system for multiple users and have suggested that the museum is the perfect place in which to produce a true physical hypermedia system. I have then pointed out some of the ways in which this might be done using existing technologies. Finally I have suggested some ways in which the museum as an environment might be made predictive and intelligent so that it can enter a truly productive partnership with its operators and visitors. I recognise that this all sounds a bit like setting up the museum as a Panopticon, all seeing, all knowing and keeping track of everybody, but really I am suggesting little more than the means by which the museum could become just like any other person, being able to see and communicate, inform, interact and collaborate with the visitor as though they were given a personal guide to help them in their tour.

References

1. Vannevar Bush “As We May Think” *Atlantic Monthly*, July, 1945.
2. Douglas B. Lenat “Artificial Intelligence” *Scientific American*, September 1995.
3. Rodney Brooks, *Cambrian Intelligence: The Early History of the new AI*. A Bradford Book, MIT Press, Cambridge, Mass, 1999.
4. Cynthia Breazeal, “Robot in Society: Friend or Appliance?”. In *Agents99* workshop on emotion-based agent architectures, Seattle, WA. August 18-26, 1999.
5. Phil Husbands, Inman Harvey, D. Cliff, and G. Miller, “Artificial Evolution: A New Path for Artificial Intelligence?” *Brain and Cognition* 34, 1997.
6. Stephen Jones “Intelligent Environments: Organisms or Objects?” to appear in *Convergence*, summer issue 2001.