James W. Strong
University of Adelaide
james@netspot.com.au

Robert F. Woobury
University of Adelaide
rw@arch.adelaide.edu.au

Psyberdesign: Designing the Cognitive Spaces of Virtual Environments

Increasingly, we find ourselves spending more time on a daily basis engaged in a variety of virtual environments, ranging from those discovered when using the stand-alone computer, to more complex distributed networks such as the World Wide Web. Virtual environments are not restricted to the popular and hyped notion of immersive Virtual Reality systems, though, in immersion, such systems provide a dimension of experience sorely lacking in most human-computer interfaces. The design of a diverse range of virtual environments, from textual through to three dimensional, would seem to require insight from those who habitually create immersive experiences, whether real or virtual. The former include architects, the latter the authors of computer games.

Virtual environments such as the Internet, and the World Wide Web in particular, are becoming increasingly confusing to navigate. Exploratory behaviour in these environments requires extensive cognitive effort, and often results in disorientation and a sense of anxiety. This paper attempts to address issues of cognitive mapping in virtual environment design, and the exciting role that architects should occupy in the creation of better virtual environments.

A virtual environment tool, called WOMBAT, has been developed to discover more about the relationship between real environment and virtual environment navigation and cognitive mapping, and consequently the degree to which concepts and theories from real environment design and cognitive mapping research can be translated to the virtual environment domain. Both a natural environment and an information environment are being investigated using WOMBAT, with the primary interest being the cognitive mapping and wayfinding activities that are exhibited during exploration.

Psyberdesign: la conception des espaces cognitives des environnements virtuels

De plus en plus, nous nous trouvons chaque jour impliqués dans une variété d'environnements virtuels, allant de ceux découverts en utilisant un ordinateur isolé, jusqu'à ceux se situant dans des résaux distribués complèxes, tels que le 'World Wide Web'. Les environnements virtuels ne sont pas limités à la notion populaire de systèmes virtuels 'immersifs', quoique, lors de l'immersion, ces systèmes offrent une dimension de l'expérience qui est malheureusement absente dans la plupart des interactions homme-ordinateur. La conception d'une gamme d'environnements virtuels, du textuel jusqu'à l'expérience trois-dimensionnelle, requiert les conseils de ceux qui habituellement créent des expériences immersives, qu'elles soient réelles ou virtuelles. Les premiers comprennent les architectes, les deuxièmes les auteurs des jeux sur ordinateurs.

En effet, on pourrait affirmer qu'une telle étude a le potentiel d'être encore plus importante dans ce nouveau domaine. Le comportement exploratoire à l'intérieur de ces environnements requiert un effort cognitif considérable, et résulte souvent en un sens d'anxiété et de désorientation. Ce papier tente d'adresser les problèmes de cartographie cognitive ('cognitive mapping') en conception d'environnements virtuels, et le rôle intéressant que les architectes devraient occuper dans la création de meilleurs environnements virtuels.

Un outil pour l'environnement virtuel, nommé WOMBAT, a été développé afin de mieux comprendre la relation entre la navigation des environnements réels et virtuels et la cartographie cognitive, et par conséquent le degré auquel les concepts et les théories provenant de la conception d'environnements réels et des recherches sur la cartographie cognitive peuvent se transférer à la conception d'environnements virtuels. Un environnement naturel, ainsi qu'un environnement virtuel sont en voie d'être examinés en utilisant WOMBAT, l'intérêt principal étant la cartographie cognitive et les activités relatives à la recherche d'un trajet exhibées au cours d'une exploration.

introduction

Kevin Lynch announced the importance of cognitive representations in relation to cities almost four decades ago in his book *The Image of the City.* He made the critical connection between psychology and design, and his intention was to encourage and assist urban planners in designing more *legible* and *imageable* cities. Unfortunately too little of Lynch's vision has manifest itself in urban design practice, and we are still confronted with confusing environments that are uncomfortable and awkward to navigate. Perhaps Lynch's advice arrived too late to save our cities, however other design challenges may benefit from Lynch's insight and lessons from psychology. Passini, building in part upon the work of Lynch, developed theory that extends to the design of individual buildings. Virtual environments desperately require a similar treatment to that which Lynch and Passini applied to cities and buildings. In fact, it could be argued that such treatment has the potential to be even more consequential in this new domain. Virtual environments such as the Internet, and the World Wide Web in particular, are becoming increasingly confusing to navigate. Exploratory behaviour in these environments requires extensive cognitive effort, and often results in disorientation and a sense of anxiety. This paper attempts to address issues of cognitive mapping in virtual environment design, and the exciting role that architects should occupy in the creation of better virtual environments.

cognitive maps and wayfinding

Cognitive mapping is the human process of developing mental representations of an environment, the result of which is a cognitive map. Such mental maps comprise not only spatial relationships, as seen in geographical maps, but also auditory, sensory and emotional data. The spatial component of a cognitive map does not have the consistency of a printed map, rather it is layered, multi-faceted, and fragmented. In almost all cases a person's representations of an environment, when expressed diagrammatically, depart quite markedly from accurate plans. Individual differences in cognitive maps occur due to varied knowledge and experience of the environment. An elaborated region on a particular person's cognitive map may be a result of frequent visits to the region, or perhaps an emotional attachment to the place. Major differences are apparent between children at various developmental stages, and between male and female adults. Cognitive maps are often shaped by exposure to existing maps such as guide maps, street maps, and plans, and as such are very important for survival in the environment. When a fully integrated and holistic cognitive map of an environment is fragmented, often "route" information plays a role in navigation. A route map is one containing data about various features in the environment that occur in a particular sequence and aid a person navigating the environment through confirmation - this occurs when items from the cognitive map are matched against corresponding features in the immediate environment. For example when driving to university from home - a very well traversed route - a holistic cognitive map usually aids navigation. Route information in such circumstances can be used to confirm location, but is rarely required for a successful trip. In a relatively unfamiliar environment, however, only fragments of information about the overall map exist. In such cases navigation, or wayfinding as it may more accurately be termed, is often guided by features in the environment, such as a distinctive corner store or bridge, which match route information stored in the person's cognitive map. Much of the information required to connect the disparate elements of incomplete cognitive maps is acquired through actively engaging in wayfinding activities, while print-based maps help to correct anomalies and bridge remaining gaps in knowledge.

from cities to cyberspaces

In his book *The Image of the City*, Lynch relates the physical characteristics of the city to the structure of the mental image. The physical characteristics are divided into a number of key elements: paths, edges, districts, nodes and landmarks. The extent to which these elements can be organised into a coherent pattern is an indication of how *legible* a city is, and a legible city is a prerequisite for a clear mental *image*. City form that has a high degree of *imagibility* assists occupants to navigate and cognitively comprehend their surrounds, and through doing so also "*heightens the potential depth and intensity of human experience*" (Lynch, 1960:5). In designing virtual environments (cyberspaces), attention should not be exclusively concentrated on aspects of efficiency, but instead distributed between promoting depth of experience, a sense of place, and creating efficient space. Lynch provides a number of design guidelines for urban planers to assist them in the task of designing the clear, rich and rewarding cities that his vision encapsulates. The approach and methodology Lynch adopted may prove to be more significant than his specific guidelines in the context of virtual environment design - attempts to directly transpose Lynch's rules to virtual environments should be approached with caution. Lynch's statement that "*we must consider not just the city as a thing in itself, but the city being perceived by its inhabitants*" (Lynch 1960:3) retains its pertinence when we consider the virtual environment instead of the city. Virtual environments are designed for people, and we should therefore focus on how people perceive and understand them.

Passini, in his book Wayfinding *in Architecture,* also presents some ideas that are highly relevant to virtual environment design. Passini is primarily concerned with wayfinding - the human process of actively exploring and navigating environments. He emphasises the experiential nature of cognitive map development and proposes that wayfinding is fundamental to our appreciation of the environment. His research is partially based upon the work of environmental psychologist, Roger Downs, who considered cognitive mapping a process rather than a product (Passini 1984:45). The process of wayfinding involves active engagement with the environment, and a continual comparison and integration of existing knowledge with stimuli from the immediate environment. In order to develop useful cognitive maps, it is therefore necessary to journey through an environment. An everyday example of this can be observed by comparing the level and quality of cognitive mapping between the drivers and passengers of cars. Passengers are likely to be relatively passive in their exploration of the environment because they are under no obligation to navigate (unless they are co-piloting). For drivers, navigation is essential and such active engagement results in a more accurate and extensive understanding of the spatial relationships in the environment

The World Wide Web is one particular virtual environment where wayfinding theory is relevant. Our current experience of using the Web is usually, if not always, characterised by a disjointed journey. In fact, metaphors such as journey and surfing, while used quite frequently in literature to describe the process of accessing Web pages, struggle to evoke a sense that resembles the experience. The notion of journey and exploration has also increasingly become absent due to the advent of powerful search engines. The common experience of the Web has become one of refining search terms rather than exploration and travel. It is proposed that the process of Web exploration, not just the structure of the Web alone, is increasingly becoming fundamental to our appreciation of the Web as a virtual environment. As the Web is considered primarily a tool for information discovery and retrieval, its appreciation is not always associated with how we experience it. However, as more of our daily life is conducted in virtual environments, imageability and wayfinding may emerge as important factors. Unfortunately a conundrum exists whereby an increase in efficiency through introduction of shortcuts like search engines, often results in a reduction in amount of time spent on important exploratory processes. A simple analogy for the search engine and Web is the catalogue and library. The catalogue bypasses the need to scan the library shelves and therefore reduces the likelihood of this often-fruitful activity. When designing virtual environments, a balance between the functional and the cognitive components must be achieved–the latter has received very little attention to date.

city/landscape metaphors

Given the theory of Lynch and Passini, it is very tempting to adopt a city or landscape metaphor when designing hypertextual virtual environments such as the Web. The following section reviews several attempts to use such metaphors, and illustrates some of the difficulties associated with their implementation.

Metaphors are undoubtedly very important in the design of computer user interfaces as they help to highlight similarities and differences between existing entities, and in doing so enable us to understand new entities through our current knowledge. The Macintosh user interface is a prime example of the huge success and importance of the metaphor.

Is a city or landscape metaphor appropriate for virtual environments, however, and if so does this enable us to apply the theory of Lynch and Passini to virtual environment design? For a virtual world that is three dimensional, devoid of any 'magic features' such as teleportals, and based on real-world entities such as mountains, buildings and streets, a landscape metaphor is clearly suitable. This type of virtual environment is a simulation of the real world. A study on wayfinding (Darken 1996) in large-scale virtual worlds concluded, by way of an experiment, that real-world environment design principles derived from the work of Lynch and others could be applied to virtual world design.

This is comforting to discover, as such conclusions imply that a substantial proportion of our existing design knowledge, including the contributions of Lynch, can be applied to the task of virtual environment design without any further deliberation. However, a virtual world as described above is only one limited type of virtual environment - what of hypertext environments such as the Web? Deiberger (1994) has attempted to explore this by using a city metaphor for textual virtual environments. Deiberger proposes the notion of an 'Information City' where buildings are the containers for documents, navigation is a process of walking the streets, catching a taxi, or taking the subway, and the whole city is structured according to an abstract distance concept based on document similarity and user needs. The whole premise for using a city metaphor is unfortunately grounded on the assumption that people rarely become lost in a city. This premise is questionable, and in addition it could be argued that becoming lost in a city is only the extreme manifestation of poor *legibility* and *imageability*. A person may be confused and uncomfortable in a city due to their incomplete and incorrect cognitive maps of the city without necessarily being lost. Such considerations are important if the city is to be used as a metaphor to specifically improve our experience of browsing hypertext environments.

Anomalies can be problematic for those who advocate the use of a city metaphor in hypertext environments. The phenomenon of "teleporting" is the most salient example. "Teleporting" is not problematic because it is an anomaly per se - rather it is problematic because of the way it is often implemented. Transition and the sense of travel are very important for the creation of cognitive maps. On the Web we lack this sense of travel– "Connecting to site…" along with a several other status messages and the reloading of the browser window, are the only indications of transition. Deiberger acknowledges the challenge that 'teleporting' presents his use of the city metaphor, and proposes that such an action occur through an underground network analogous to a subway. This transport mechanism is instantaneous–there is no sense of travel. The adoption of a subway metaphor is rather perplexing given that the subway is arguably the principal contributor to confusion and disorientation when navigating large cities[2]. Deiberger dismisses an alternative metaphor of "flying" as *"[it] is too distracting and visually intense to be useful for general hypertext linking"* (Deiberger 1994) This clearly highlights the dilemmas that a highly direct implementation of the city metaphor faces. The metaphor becomes the dominant visual element and distracts the user from the real purpose of being in the virtual environment. Perhaps this is partly the reason for such a slow uptake of VRML (Virtual Reality Markup Language) by Web users, despite the expectations of those close to the technology.

When considering the use of a city metaphor, or any metaphor, we need to "*situate [ourselves] in the problem domain rather than immediately fixating on a particular metaphor*" (Coyne 1995:252). The problem domain of virtual environments is largely cognitive, and is one of communicating the relationships between information. Just as Lynch was interested in "*the city being perceived by its inhabitants",* we must concentrate on the virtual environment as it is perceived by its inhabitants. Lynch and Passini partly based their design principles and solutions to problems of *legibility*, *imageability*, and wayfinding on psychological theory. We need to understand how people perceive and cognitively map virtual environments. The task is arguably more complex than understanding the relationship between the real world environment and cognitive mapping, however a similar approach is required. The landscape metaphor is useful in that it hints at notions of travel and discovery, which are important factors for the creation of cognitive maps, however is should not necessarily be the primary tool.

designing cognitive maps

In a sense, architects and urban planners are already designing cognitive maps. An extreme example of how architects can manipulate or induce a particular organization within a cognitive map is in the design of labyrinths or maze-like structures. A consideration of cognitive mapping is not usually an explicit part of designing physical buildings or urban environments, but when considering virtual environments the cognitive component of design is essential. Virtual environments are not geographic spaces–they are cognitive and social spaces. Architects will still have a contribution to make in the design of visual elements, however their role is not one of trying to make the virtual environment something visual, but of helping to shape the cognitive maps of its inhabitants. The design principles that these architects use will be derived partly from scientific research into the relationship between virtual environments and cognitive mapping.

We know how people physically move in cities–either by walking, driving, taking a taxi, or the subway. The possibilities are limited, and the properties of each mode are somewhat fixed–for example, speed in walking varies little between individuals. However, in the virtual environment we do not understand nearly as precisely the processes by which people move or experience the environment. In Virtual Reality the environment may be experienced by walking, running or flying, but across different VR systems these processes are likely to be quite different in nature – there are no universal standards. However, more significantly we know very little about how people explore and understand virtual environments. In fact, our understanding of wayfinding in the real world is still at a relatively immature stage. In order to solve the problems inherent in current virtual environment design, we need to better understand the various modes of travel as well as wayfinding and other cognitive activities.

mapping the web

There have been many attempts to visually represent parts of the Web in order to assist people in information discovery and retrieval. However, such attempts have not tended to concentrate on the way people cognitively represent such environments. Rather they have focussed on providing access to the Web through external metaphors such as landscapes, connected networks of rooms or the ubiquitous desktop. In contrast, the cognitive maps people generate have concepts in common, such as routes, landmarks and boundaries, yet are idiosyncratic in the composition of these conceptual objects, and are actively and continuously updated with experience in an environment. Even Web search engines, seemingly the antithesis of a cognitive mapping approach, are now providing features that cluster together related concepts in what might be considered a dynamically updated map of that part of the Web they uncover. For example, the Lycos search engine provides categories for Web site reviews, pictures, personal home pages, a bookstore search, yellow pages and news search for any specified search.

Our thesis is that the cognitive maps people generate are especially important when considering complex computer networked environments such as the Web. Providing people with a map of an environment, in either a physical or virtual context, is likely to improve their ability to understand the relationship between elements. However, having a map is only one side of the cognitive mapping coin–cognitive maps are created and updated through experience in an environment. In effect, a cognitive map records the process of wayfinding, which is an active and continuous process involving the integration of information from stimuli in the immediate environment with existing cognitive representations.

One example of such an attempt to map the Web, and in particular Web sites, is HyperSpace. This system organises the information not according to geographical location, but according to a user-defined structure, so that related topics are displayed adjacent to each other, and unrelated topics are spatially separated (Figure 1.).

Another tool, which caught the brief attention of some Web users, is Apple Computer's HotSauce (Figure 2). HotSauce allows users to fly through

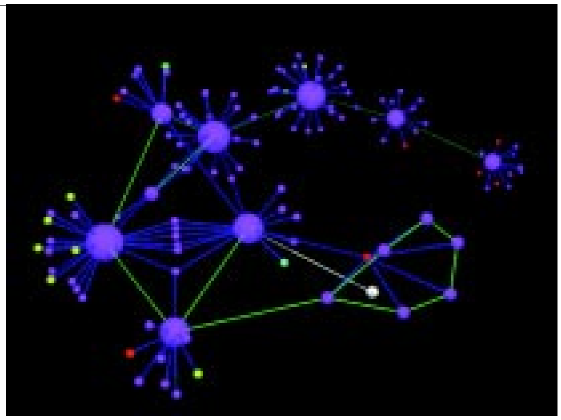

Figure 1. The hyperspace system allows a user to create a realtime visualization of the structure of a set of web pages while browsing through them. See http: / / www.cs.bam.ac.uk / ~amw / hyperspace / .

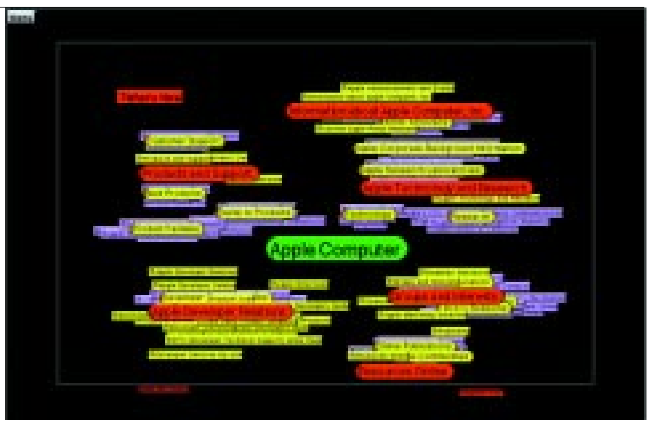

Figure 2. HotSauce from Apple Computer allows users to fly through layers of labels representing hypertext links.

layers of labels, with each label linking to a Web page. Interestingly, HotSauce is no longer available or supported by Apple. The provision of maps is certainly important in that it assists people to navigate and develop clear cognitive maps, however the overarching process of navigation through the Web with these aids does not seem to be considered in such implementations.

Recently there has been a proliferation of programs (e.g. Geoboy) that trace paths from a person's computer to a Web site. These paths are based on guessed geographical locations of Internet gateway machines (routers). "Traceroute" tools provide extra information such as the ping times (lag) between the various nodes along a route. It is interesting that such tools have been so widely adopted by Web users, as they seem to be regressive in their attempts to associate the physical with the functional configuration of the Web. Perhaps the acceptance of such tools is an indication that people are in dire need of extra cues and information as a basis for generating clear cognitive maps. Experienced Web users often take note of the URL (Uniform Resource Locator) in the location box of their browsers to gain extra information about the resource that they are accessing. The URL can provide information about the institution and country in which the resource is located. URLs often contain information about the directory structure of a remote Web site, which may also be used in generating cognitive maps. There are possibly many other less explicit cues employed by Web users to assist them in cognitive map creation.

Lynch would have loved VE's

At this point the reader may have the impression that all virtual environments are cognitively confusing and uncomfortable to use. Many people, however, claim they experience no significant problems when navigating virtual environments such as the Web, even though they occasionally find them confusing and frustrating. People do complain about elements of the content such as 'flashing' images, the use of frames, scrolling text or too much text, as well as technical issues relating to the network such as 'lag' and bandwidth congestion. Lynch, referring to the citizens of American cities says:

> *"They are clear enough about the ugliness of the world they live in, and they are quite vocal about the dirt, the smoke, the heat, and the congestion…but they are hardly aware of the potential value of harmonious surroundings . . . they can have little sense of what a setting can mean in terms of daily delight, or as a continuous anchor for their lives, or as an extension of the meaningfulness and richness of the world"*" (Lynch 1960:2).

Citizens of virtual environments face a similar diagnosis. Lynch was keen to improve the quality of cities, not by addressing issues such as the smoke, heat and congestion, but by improving an often hidden element relating to the overall structure of the city as perceived by its citizens. To this end Lynch would have reveled in the task of virtual environment design. Firstly, experimentation is more readily performed. Virtual environments, by their very nature, allow us to capture vast quantities of information about the activities of their citizens. Experimentation is much cheaper, and therefore allows for larger and more representative samples of the population to participate. Lynch regretfully ignored individual differences due to the complexity it would introduce, but in the virtual environment such factors can more readily be considered. By understanding cognitive mapping more accurately, and through examination of individual differences, design rules can be applied with more certainty. There is no need to rely on a public image to inform the invention of design rules for virtual environments.

Secondly, the application of any emerging design rules to existing virtual environments is not likely to be nearly as labor intensive, thus enabling validity to be directly tested. In the physical world, the scope for dramatically changing the environment is obviously restricted. Finally, the value of cognitive mapping and *imageability* of the environment is possibly more important in the virtual domain, than the physical domain.

Hopefully this provides some optimism for virtual environment designers. The problems of cognitive mapping in virtual environment design may be less visible than other more superficial problems, but is a very real concern. While the virtual environment presents many challenges for designers, it is important to realise that it also allows for a greater level of experimentation.

experimental mapping

An essential prerequisite for study of cognitive mapping in virtual environments is a virtual environment. To be applicable, such an environment should have a conceptual structure in common with the Web, namely a network of interconnected objects. To be accessible to our senses, both the nodes and the connections should be capable of being experienced. To allow the cognitive mapping theory of Lynch and others to be applied in information worlds as well as physical worlds the environment should be able to be interpreted in both contexts.

We have built WOMBAT, a virtual environment tool with some of these characteristics. Abstractly, WOMBAT comprises the following:

- a network of bi-directionally interconnected nodes;
- an ability for both nodes and their connections to carry information;
- a user interface that displays nodes and the traversal of connections, locates a user within a network and provides a set of navigation commands.

A native Australian bush landscape[3] provided the incubator for WOMBAT in the sense of being the real environment which WOMBAT first modelled. Like its real-life cousin (a slightly dim-witted heavy and slow marsupial) WOMBAT works by moving along paths in a landscape. The information available to a WOMBAT user is what can be sensed along a path. For the bush environment WOMBAT employs segments of real video footage of the actual bush landscape to simulate movement down a path as well as the process of "panning" from path to path at nodes (the intersection of paths).

In order to provide a consistent and extensible structure for the video segments, a two-manifold surface with boundary was chosen to create

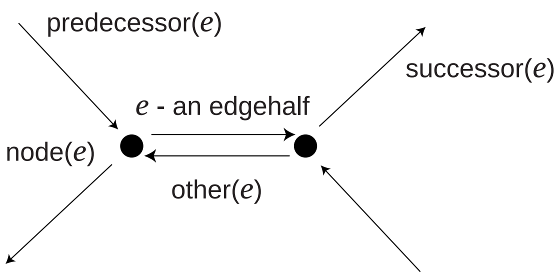

Figure 3. The split edge data structure. Associated edgehalves are the functions predessor (returning the precessor edgehalf on the face of an edgehalf), successor (returning the successor edgehalf on the face of an edgehalf), other (returning the edgehalf adjacent to an edgehalf on an adjacent face) and node (returning the node from which the edgehalf emanates).

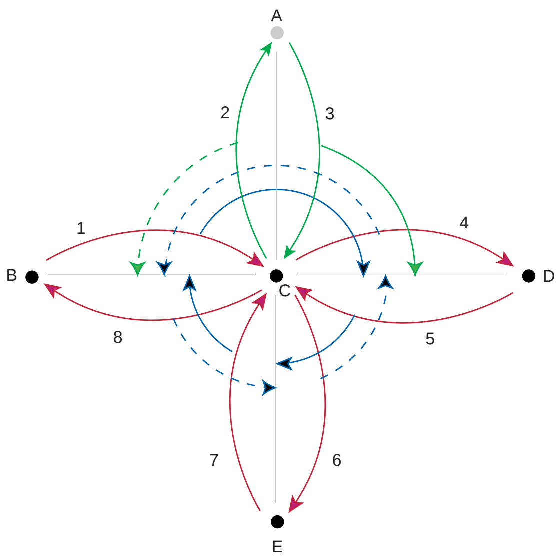

Figure 4. Illustration of "stub" edgehalves. A T-intersection exists where B-C-D intersects C-E. Normally edgehalves are associated with a segment of video travelling between two nodes (for example edgehalf 4 contains the video segment between nodes C and D). Each edgehalf is also associated with two pans (only one pan is depicted in the diagram as only node C pans have been considered). When a user arrives at C after coming from D, pans must be obtained from "stub" edgehalves (in this case edgehalves 2 and 3, indicated by green lines). These two edgehalves have no video associated with them–there is no way of travelling from C to A.

WOMBAT's network. The operations that can be performed on such a surface are consistent with the requirements for a virtual environment tool for exploring cognitive mapping and wayfinding. For an information environment, the structure allows for bi-directional links which may be essential in providing more structure and coherence to the navigational experience. Two-manifold surfaces form a basic structure in solid modelling theory (Mantyla 1988). Informally, a two manifold is a continuous surface–every real object has such a surface as its boundary. Unlike the surface bounding a real object, in which every point on the surface has other points on the surface around it, our surfaces themselves have boundaries. Solid modelling theory shows that two-manifold surfaces can be modelled in a variety of ways, by recording relations between nodes, edges and faces (parts of the surface bounded by cycles of edges). Figure 3 shows one such scheme, which WOMBAT uses, the split-edge structure in which each edge is associated with two edgehalves, each being part of a cycle of edges about a face adjacent to the edge.

In WOMBAT each edgehalf carries a video segment depicting the process of moving along that edgehalf through the bush. Each edgehalf also carries the two video segments, panning to the predecessor and successor edgehalves respectively. Due to technical reasons, video pans cannot be concatenated. Each pan must be visited manually in order to rotate 360 degrees. An anomaly arises when considering a T-intersection, as no edgehalves exist to carry the two necessary pans. The solution is the addition of "stub" edgehalves, which carry pan segments, but not video segments depicting the process of moving along that edgehalf. These "stub" edgehalves are entirely consistent with the two-manifold conditions and are modeled in exactly the same data structure as the rest of the network (Figure 4). The natural operators over boundary representations are the so-called Euler operators, which make changes in a data structure such that the structure remains consistent as a model of a two-manifold surface after every change. WOMBAT uses a set of such operators to construct its networks in order to ensure consistency and easy maintenance as the environment expands.

WOMBAT provides a simple user interface (Figure 5), allowing a single user to move along paths and pan at nodes. In the bush environment these two acts access all the node and path information available. In another type of information environment additional user interface operations might be required to experience node and path information.

WOMBAT has some similarities to a computer adventure game, such as Riven or Myst, but there is no game element to the system, and WOMBAT allows for accurate recording of user activity, and the ability to easily add extra functionality for experimental purposes. WOMBAT's few users have noted the same sense of adventure and exploratory behaviour elicited by computer games. WOMBAT uses standard Web technology for its interface and can be used in an intranet environment.

The environment chosen to raise WOMBAT is a partly native bush setting in the Adelaide Hills, which contains numerous human-made pathways. The terrain is undulating and contains patches of dense bush, an open field and a lake. The environment was chosen because it readily evokes a degree of confusion or surprise in anyone who chooses to navigate through it. The first author (Strong) has lived in the area for more than 20 years and still finds aspects of his cognitive maps incomplete or entirely wrong! A natural environment, as opposed to man-made environment, was chosen because of the lack of a predefined grid. The goal of experimentation at this stage is not to necessarily confirm that Lynch and Passini's theory is valid in a virtual environment, rather to explore cognitive map generation and wayfinding behaviour in a virtual environment using a similar methodology. A city environment, either real or computer-generated, could be used to more directly test Lynch's urban design guidelines.

nodes on a notional grid

The nodes are derived from a notional 20mx20m grid placed over an aerial photo of the environment. Where horizontal and vertical lines between nodes intersect existing well-defined pathways, extra nodes are generated. (Figure 6).

Figure 5. The WOMBAT interface. Users can either proceed straight ahead or rotate left or right. Each operation will display a segment of video.

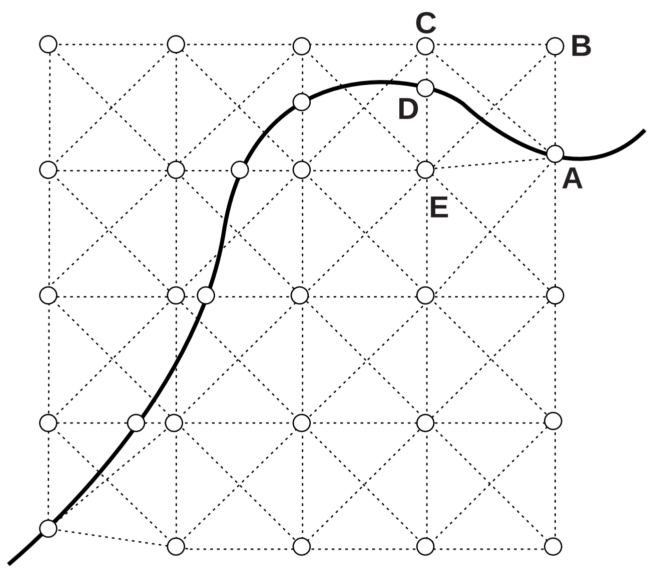

Figure 6. How nodes are generated. The distance between C and B is 20m. Note how a node is created at D by the intersection of the line between C and E. Node A is has been shifted up a small amount to the location of the node on the pathway created by an original A-B intersection. In cases where the nodes are very close together (less than 5m) the node is moved to the location of the intersection.

Segments of video footage were filmed between the nodes in both directions. These segments were typically 5-15 seconds in length. The video footage between nodes was sped up by a factor of two to reduce the time it takes to travel between locations. A 360-degree pan at each node was also captured. The footage was taken while on foot, and the result is quite "jerky," however this was considered acceptable for experimental purposes. The video was transferred to a computer and converted to QuickTime format.

WOMBAT adopts one aspect of the landscape metaphor, namely, that information is located on and accessed by movement on a continuous surface. Unlike other landscape metaphor systems it immerses a user in this surface, deliberately not providing an ability to "pop-up" above the surface to view it from the "outside." This feature arose from WOMBAT's inception as a virtual environment for both real and virtual places, and its first application to a real place in which movement along a surface is dominant. The solid modelling theory, which makes the large collection of path traversal information coherent, depends upon having a surface to model, even though this surface may be arbitrarily shaped.

using WOMBAT

Several people have experienced WOMBAT in its native bush setting. Most agree that it presents a convincing simulation of movement in that landscape and that, by using it, a mental map of its modeled area could be formed. We are embarking on a comparative use of WOMBAT, in both the bush and an information environment, towards the end of learning more about what concepts from navigation in real environments translates to the virtual and as a test-bed for new navigation, mapping and wayfinding devices. In this section we discuss our current work with WOMBAT.

A number of participants will be recruited to take part in the experiment. Data about a participant's exact movements in the environment will be recorded by the system. All participants will be questioned on their navigational experience during and after exploration of the environment–did they experience a sense of place anywhere in the environment, and did they encounter any confusion at any stage during their exploration? The participants will then be asked to graphically represent their cognitive maps, as 2D drawings or simple 3D computer models.

In the bush environment, the use of video between nodes (traversals) and at nodes (pans) provides a definite sense of travel. However, if the segments of video between nodes are replaced by still images, resulting in a more disjointed and hypertext-like environment, would this sense of travel be lost, and what differences, if any, may exist in participant's representations of their cognitive maps? The substitution of still images for video segments will result in an experience more akin to 'teleporting' than travelling. Those familiar with multi-segment QuickTime VR movies will be able to identify with this phenomenon. Teleportals and other short cuts and navigational aids may be introduced into the environment and the resulting cognitive maps compared with those obtained from exploration of more realistic environments containing fewer anomalies.

Comparison of wayfinding and cognitive mapping in WOMBAT, with similar activity in the real bush environment will also be made. A small proportion of the participants for the real bush environment will have also used WOMBAT, so that feedback can be gathered regarding any perceived differences with specific regard to wayfinding and cognitive mapping. A close correlation would enable WOMBAT to be used for analysis of existing real world environments to better understand wayfinding and cognitive mapping in buildings, cities, and natural environments.

Additional navigational aids will be provided in an experimental fashion. For example an aerial photograph or plan with a 'you are here' marker may be activated for certain regions of the environment, or only made available to a subset of participants. In addition, participants' paths through the environment may be traced on the plan to assist in comprehension of the environment. We would expect provision of such aids to significantly increase the coherence of cognitive maps, and therefore comprehension of the environment. WOMBAT users will also be able to mark locations in the environment and return to them from any node–a function equivalent to "bookmarking" on the Web.

WOMBAT provides some exciting possibilities for an information environment. The value of bi-directional links is particularly of interest. IBM's Aqui project explores the potential of such bi-directional links (see http://www.aqui.ibm.com). The Web is currently uni-directional–a link from document "a" to document "b" does not automatically imply that any link is or should be expressed between document "b" and document "a." The "back" button on browsers partially addresses this inadequacy, however it is only session-based and problems arise when multiple paths are introduced. True bi-directionality allows a user to travel down, as well as up, hypertext paths. Where there are multiple paths to a destination, a fork will exist that allows the user to select which path to follow. Another important benefit of bi-directional links is the ability to determine all paths from the user's current location to an arbitrary destination. This is useful in the case of bookmarking–when a user chooses to return to a bookmark, the system can present an array of routes to the bookmark, thus reinforcing previously visited links. WOMBAT will create a map of paths as the user explores an environment. By providing path information, WOMBAT introduces an important component lacking from the current Web browsing experience.

The ability to add information in the links between documents provides further experimental posibilities. Each node can carry a hypertext document, and each edgehalf between two nodes can carry meta-data. This meta-data could provide concepts connecting two documents, thereby providing a more continuous navigational experience.

As for the bush environment, navigational aids will be provided in a similar experimental fashion, and participants will be asked to construct diagrammatic representations of their cognitive maps. By reinforcing the path and instilling a sense of travel in its users, WOMBAT may provide participants with more substantial and effective cognitive maps. The tool will assist in the discovery of elements from natural and urban environments that may be utilized in virtual environments such as the Web.

conclusions

The process of looking at cognitive maps has been undervalued in virtual environment design. The creation of WOMBAT has strengthened convictions that cognitive mapping and wayfinding theory, as promoted by Lynch and Passini, is applicable in the design of virtual environments at an abstract, not just visual, level. Browsing the Web is becoming more disjointed and exploratory behaviour is vanishing as powerful search engines become the primary mode of navigation. This is having an adverse impact on cognitive map generation, and consequently our appreciation of the Web, because transition and a sense of travel are especially important for the development of useful and coherent cognitive maps.

endnotes

[1]The term *virtual environment* encompasses, in this paper, phenomena such as Cyberspace, Virtual Reality, Internet, World Wide Web, and MUDs (Multi User Domains / Dimensions / Dungeons). Emphasis is particularly on the networked nature of these phenomena not the visual aspect, thus extending the notion of a virtual environment beyond three dimensional Virtual Reality. For example, a University Department's local area network of Macintosh computers may be considered a virtual environment.

[2] Passini uses Montreal's metro in some of his experiments and notes the difficulties that his subjects encounter when trying to formulate clear cognitive maps.

[3] "Bush"= "forest" or "woodland" for readers from "up-over."

references

Coyne, R.D., 1995. *Designing Information Technology in the Postmodern Age.* Cambridge, MA: The MIT Press.

Darken, R., 1996. Wayfinding *in Large-Scale Virtual Worlds.* Dissertation. The George Washington University.

Deiberger, Andreas, 1994. *Navigation in Textual Virtual Environments using a City Metaphor.* http: / / www.lcc.gatech.edu / ~dieberger / Thesis / Thesis.html (April 15, 1998).

Gaver, William, 1995. *Oh What a Tangled Web We Weave: Metaphor and Mapping in Graphical Interfaces.* http: / / www.acm.org / sigchi / chi95 / Electronic / documnts / shortppr / wwg2bdy.htm (April 15, 1998).

Lynch, Kevin, 1960. *The Image of the City.* Cambridge, MA: The MIT Press.

Mantyla, M., 1988. *An Introduction to Solid Modeling.* Rockville, MD: Computer Science Press.

Passini, Romedi, 1984. *Wayfinding in Architecture.* New York, NY: Van Nostrand Reinhold.

Ullmer, Brygg, 1995. *Towards "Virtual Reality": Explorations in Spatializing Web Content* http: / / ullmer.www.media.mit.edu / people / ullmer / papers / urbcyber / (April 15, 1998).