Spaces, Places, Landscapes and Views: experiential design of shared information spaces

John A. Waterworth

Department of Informatics
Umeå University
S-901 87 UMEÅ
Sweden
+46 90 16 67 31
jwworth@informatik.umu.se

ABSTRACT

This chapter focuses on the World Wide Web (Web) as a provider of shared information landscapes. It reviews our work to design 3D spaces for information navigation and social interaction, and suggests an approach to such design based on an experiential theory of meaning. The increasing use of virtual 3D space in information environments is noted, and Personal Spaces are contrasted with Public Places. Earlier work on Information Islands, Vehicles and customisable Views of such information spaces is also presented. The experiential approach, as applied to information landscape design, is contrasted with the traditional view of Human-Computer Interaction (HCI) design as a means of conveying system functionality from the mind of the designer to that of the user. This experiential approach seems promising, if we assume that we do not know in advance what the functions of interactions in shared information spaces might be. As with life in general, such interactions mean what they are experienced to be.

1 Introduction

There are many ways in which the world's most popular hypermedia system (by far), the World Wide Web (Web), does not reflect the hypermedia usability research that preceded it (see [1] for a catalogue of what were considered the key research issues at that time). Perhaps the most unexpected thing about the Web as a whole is that no-one is designing it. Three other ways in which it has not conformed to what was expected of hypermedia, are: first, the use of three-dimensional graphics to give a sense of space; second, the fact that there is one Web which all users cohabit; and, third, the fact that we can communicate with each other from within the Web.

This chapter develops three themes which follow from these unexpected characteristics of the Web: Personal Spaces versus Public Places, the notion of Vehicles with Views, and the potential for Presences and Concealment. These are illustrated with some recent examples of our work, which adopts an experiential approach, which contrasts with the traditional view of HCI design as facilitating the communication of functionality between designer and user. This work has been motivated by the realisation that a profound change is taking place with the evolution of information and communications systems into self-contained virtual environments. Inhabiting a virtual world is very different from using electronic tools in the real world: hence, the importance of virtual presence.

2 Personal Spaces versus Public Places

An increasingly popular approach to the representation of information on the Web is to use 3D rendering techniques to convey a sense of space and apparently solid structure. This means that information explorers can bring their innate skills for spatial navigation into play, in addition to those few sensori-motor abilities utilised by the familiar direct manipulation (WIMP) interface. However, because no-one is designing the Web, and because of the simple linking mechanisms underlying its evolution, there is no way to make sense of its structure as a whole. There can be no 3D representation capturing its whole structure which, as is implied by our approach, means that people simply cannot make sense of its structure as a whole.

However, space is powerful as a means of representing the structure of designed environments, such as personal file systems and the intranets of organisations. Personal environments can be happily represented as Personal Spaces - 3D structures apparently containing stored and current items of interest to the individual user (e.g. StackSpace; [2]). Figure1 gives an indication of the StackSpace environment.

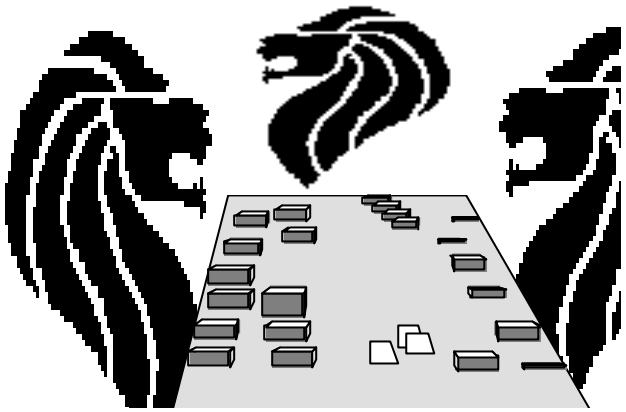

Figure 1 - A Space in StackSpace

The need for multi-threadedness, chronology and currency-tracking are taken care of in StackSpace in the following ways. Multiple stacks develop as the user explores, and the top slices are the most recent or current. The items that are furthest in (away from the viewer) are of least interest or relevance to the current task. Items decline gracefully in interest, by moving away from the viewer, but then fall over a cliff (although the retentive user could, in principle, extend the horizon to infinity). Contextual bridges (shown in Figure 2) and cords show relationships between items in different stacks. For example, the bridges shown in Figure 2 indicate that stacks A and B, B and C, and C and D all have at least one "slice" in common - the same Web page has turned up on both stacks.

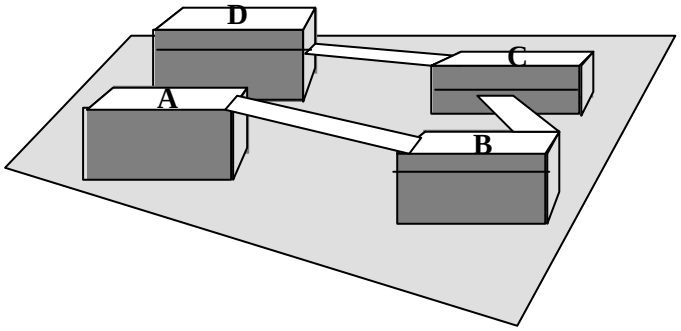

Figure 2 - Contextual Bridges in StackSpace

Computer searches, for example the results of an agent carrying out a collecting task on a topic, also produce stacks. Users then manipulate, edit and label spaces, stacks and slices to customise the material found. Edited stacks comprise views of a topic, hot-lists and traveller's tales reflecting experiences navigating in pursuit of a particular purpose. They are exchangeable with others and thus provide a means of communicating about the process and results of information navigation.

We can distinguish the idea of space from that of place. In a sense, everywhere on the Web is currently a public place (some have restricted access, but I will disregard that for present purposes), even the humble single-screen individual-user home page. They are public because anyone can go there and, often, several people will be there at the same time. But they are not aware of each other. In this sense cyberspace is unlike reality. We cannot generally see where people are in cyberspace. We have public places and personal spaces but no public spaces, because users do not share a sense of each other's presence in those 3D structures (Here I am ignoring a few emerging social spaces, such as The Palace, specifically designed for some kind of social interaction via the Web). We need to use 3D to convey aspects of human presence, to represent the people, not just aspects of the available information; but 3D alone will not be sufficient.

My use of the words "space" and "place" here is differs from that of Harrison and Dourish [3]. By my usage, "Public spaces" are shared 3D places where we can and do interact with others in real time, which reflects normal use of this term. I recognise that socialising in places implies conventions of appropriate behaviour (behaving in or out of place), a topic thoroughly addressed by Meyrowitz [4], and that social places need not be spatial. However, and contrary to Harrison and Dourish, "a place" in normal usage does not of itself imply what they mean by "placeness". A place can simply be a location. To avoid this confusion, we need to be explicit about whether we mean by "place" a particular location, or a social occasion to which a set of conventions applies. By "place" I mean the former, and "public places" are spatial or other locations open to everyone. A "Public space" is then a spatial place, which may or may not be used for social occasions to which a set of conventions applies. To use the term "place" to mean the latter seems to me confusing, since we are really talking about appropriate behaviour for a social occasion, wherever it takes place. The expression "behaving out of place" refers metaphorically to a time when different social occasions took place in different places (in the sense of location in space). We cannot, I think, design "placeness" in Harrison and Dourish's sense, we can only design locations which may or may not be spatial. Real or virtual space may reduce the tendency to behave "out of place" (as suggested by Myrowitz's thesis); the absence of such space in newsgroups may account for the frequency of socially inappropriate behaviour there - despite their being "a place".

The Web provides a marvellous medium for information exchange, for contacting others, for sharing opinions, for finding out about events, and for keeping in touch with recent developments. But as we explore the Web, we stay "at home". People can send us messages, can search for

things posted with our names attached. Once they have our address they can write to us. Maybe they can send e-mail from a page of ours they came across. But they don't know where we are at any given moment, they know only the address we use for sending information (and not always that) and the information we make available. If we have a camera set up in the office and linked to the Web, they can see when we are in our office. But they probably already knew where we worked. If they see us at the terminal we might well be navigating around the Web, but where? Bodily presence is no longer as important as where our attention is located.

We don't always want to have to go and look for things ourselves, and search engines of one kind or another are increasingly used to locate information on the Web, especially by more experienced users. The notion of software 'agents' (also known less misleadingly as personal digital assistants), which can carry out tasks for us in the background while we get on with other things, is much talked about and complimentary to the idea of using space. Agents provide services for their "masters", but the real agents in cyberspace - the people - remain unrepresented. Because of this, we cannot search for people, only for the things they have left in cyberspace. It should also be possible to enter the attributes of people we might want to locate and have the system report back where they are in cyberspace, where they have been recently, and so on, adding value to everyday reality. Some of the things we might wish to track about people include when and where social groups form, the navigational paths of individuals (or their agents), and their interests (which their agents would know about). Agents meeting with other agents could provide some of this information. Whether this is seen as threatening or not depends on the level of confidence one has in one's agents.

In summary, public 3D places are proliferating but they suffer from the same limitations as 2D places - they are inflexible yet changing. There is no space between public 3D Places, because there is no context between sites. As the Web is today, this is an insoluble problem. Personal spaces are a promising way to make sense of material gathered from the Web. They also provide mechanisms for editing and sharing materials. Public places are shared, but public spaces need ***3Dspace*** and ***visible*** people.

3 Information Cities, Islands, Vehicles and Views

About 10 years ago, the idea of a virtual 3D Information City - a way of presenting sets of information to tap people's skills in urban navigation, was raised in Singapore, itself a highly

"wired" city aiming to deal largely in information in the future. This idea was circulated in internal research reports in 1988/89, and published in brief form in [5].

An 'Information City' was also suggested early on by Dieberger [6], who has since published many papers on a textually-described 3D city (see Dieberger, this volume). Navigating a space that exists only as a text description, as in early networked "adventure" games, is an interesting task to study, especially when it is carried out as a social activity. But interpreting text is a very different skill from navigating a 3D structure, whether real or virtual. It is hard to imagine that one could use textual descriptions to create spaces that could be used for purposes other than their own exploration, since that in itself will be very demanding of cognitive resources. The point of the virtual 3D structure is to remove the need for people to use linguistic interpretation to make sense of interfaces and to tap their largely unconscious sensori-motor skills, thus freeing cognitive capacity for tasks other than navigation.

A later development of our basic Info City idea was the "Information Islands" model for the Singapore National Computer Board's National Information Infrastructure Project; this work was carried out in 1993/94 and aspects of the model were published soon after (see [7], [8], [9]). Again, this was a "natural" idea to arise on an "Intelligent Island" in close proximity to the giant archipelago of islands that is Indonesia. This section of the chapter summarises the Information Islands model, to set the central idea of Views in its original context.

Under the "Information Islands" model, the world (through which the structure of a set of information is represented) is seen as a group of Archipelagos, each composed of Information Islands. Each Archipelago represents a set of broadly related entities, providing a clear, top-level classification of what is available in this world and where it is to be found - an overall orientation that is easily accessible to both the novice and the experienced user. Each major class of service or application exists as an Archipelago. Examples might be Entertainments, Government Services, Information Services, Communications, Medical, and Financial Services. Archipelagos are collections of Information Islands. The size of an Archipelago depends on the number (and size) of the Islands of which it is composed.

Each Island generally contains only one subclass of service. Users will become familiar with this world mostly by learning the location of Islands with the kinds of services they use or are interested in. Each Island contains one or more Buildings. Some Islands may be representations of the services offered by particular providers - Provider Islands. An example might be a particular information provider’s Island located near other Information Services Islands.

Each Building contains a set of information sources or services related to a particular topic or application focus. Examples might be Weather Building, Sports Building, Stocks and Shares Building. Buildings on a particular Island will have distinctive appearance (shape, colour, graphics, text). All Buildings have common features including a Store Directory and an Information Counter (see Figure 3). The Store Directory allows users to browse and select from what is available in a Building. The Information Counter is a public agent that searches for information in response to requests from users. Buildings contain standard features to assist in navigation and item location (cf. [10]).

Archipelagos, Islands, and Buildings become bigger the more items they contain. Each Archipelago is formed by placing a boundary around the Islands from which it is composed. Each Archipelago has a distinctive colour that provides a context and reminder to the user of the focus he has chosen. As the user zooms in for more detail, the view of Archipelagos is replaced by a view of the Islands from which the selected or central Archipelago is composed. Intermediary views provide realism and orientation as the user zooms down.

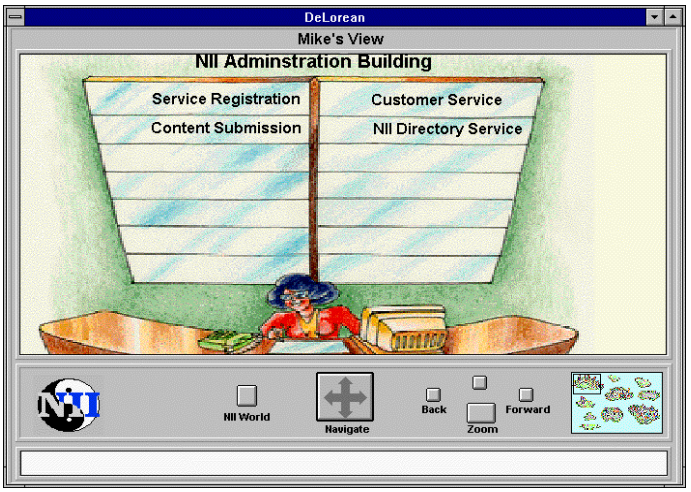

Figure 3 - The Store Directory

When a single Archipelago is shown, the Islands from which it is composed are represented separately. A view of a single Island is a map of the collections of services provided and which are represented as Buildings (see Figure 4). Buildings that are related are clustered together into no more than ten villages. Each Building contains no more than 20 Floors, and each Floor generally contains a set of related services.

The user views the available services by zooming down and selecting a particular Building, which is of a different colour from that of its neighbours. He enters the foyer (the background retains the colour of the Building to which it belongs) and can then either browse the Store Directory or

consult the public search agent at the Information Desk. The Store Directory presents a list of the service types available on each Floor of the Building. At each Floor there is a 'lift lobby' where users consult a Floor Directory (like the Store Directory, but listing individual services), to invoke the service they require.

An important part of interacting with this world of information is the exploration, selection and collection of items of interest to the individual user. These items may be services, information or particular configurations of applications. One common way of catering for this need for a personal selection from a public world (a set of public places) is to demarcate part of the world as personal, and allow the user to collect items and configure that private area. This is one of the key ideas behind the Rooms concept [11]. However, such an approach is limiting. Users must navigate to their own area frequently, bringing back items they want to collect, then venture out again into the world-at-large. In such a case, the disadvantages of a spatial metaphor can outweigh the advantages: because the users' personal space is part of the global information space, they frequently have to move around to switch between their own perspective and the higher levels of organisation. Use (which always involves a user) is confounded with level of structural organisation (which includes a User level). Use should be possible at any level, at any time. A private area at a particular location in the informational world may not be the best way of supporting individual customisation.

Figure 4 - Islands and Buildings

3.1 Vehicles with Views

To overcome these problems, the concept of private Vehicles was developed; these can be thought of as transparent, mobile, personal workspaces. They combine the idea of a private collection of information and configuration of services (customised workspace) with that of multi-

level navigational device and customised information viewer. The user is always in his (or her) Vehicle, and therefore always has access to both public and private worlds. Items can be transferred between these two without navigating space. A key aspect of the model is that the user has a filtered way of looking at the same spatially-arranged world that occupies public space. The private "world" is actually a manipulable way of viewing rather than a specific place (cf. [12]). It assumes that there is no one true view of the world, but always many possible ways of looking at things.

In the original Information Islands model, the user in his Vehicle had two Views of the world outside - a (somewhat ironically named) public “God’s Eye” View that includes everything that is available, and a personal View showing only those items that the user has selected as of interest or use (see Figure 5). He has only one set of navigation and viewing controls; the user chooses upon which View or Views they act. Although there are two Views, there is only one world. The private View and the God’s Eye View are different perspectives on the same world; the former is filtered and limited, the latter is a complete display at the level of detail on which it is focused. The user can choose to have a split screen showing both Views simultaneously, or alternate between the two. Views have some similarities with the idea of 'Magic Lenses' [13]. However, a key aspect of Views is that the 3D structural integrity of the world is always maintained (the philosophy of "one world, multiple views").

The user can also ‘yoke’ the two Views together so that the public View and the private View are then both from the same viewpoint (viewpoint: literally, the position in virtual space from which views are taken), changing together as the user navigates or inspects information at different levels. This can be useful when he wants to know what else is available at a place, other than the things he has already chosen to include in his View. This is also useful during customisation, when he can fly around the world-at-large and select things that he will then see included in his own View. At other times, the two views are 'unyoked' and he will select one or other of the two Views to be updated as his Vehicle moves, but not both. The View that is selected (private or public) will be the one that is affected by the navigation controls, the other will remain focused on where it was when last selected. He can use the public View as a navigational overview while exploring in detail with the private View. Alternatively, he can have his private View as an overview and move around the world via the public View collecting items to add to his private world. Selecting ‘yoke’ will cause the less-recently-selected View to be updated to match that of

the more-recently-selected View. This means that navigation can be done on either View, and the other View aligned to that viewpoint when required.

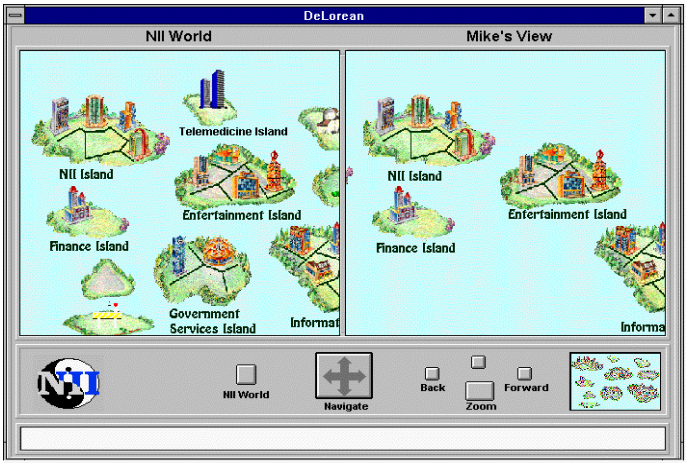

Figure 5 - Two Views of the same part of the world

The provision of both a personal and a comprehensive public View means that the user has access to a customised world, as well as the world-at-large. This customised world is a subset of the world-at-large, selected by the user but retaining the layout and grouping inherited from the larger world. Items are simply dragged from the public View to the private View window. Apart from this simplified View, the user may want instant access to a few frequently used services and applications. Two mechanisms are provided for this: the Vehicle's Memory and the Glove Compartment. The Vehicle's Memory is a list of places the user wants the Vehicle to remember, so that they can be rapidly revisited without the need for navigation. This is essentially the same as what became known as "bookmarks". When he is at a particular location, at whatever level in the hierarchy, a user may select the 'Memorise' option, resulting in that location being added to the memory list. At a later time, he simple clicks on that item in the list to instantly move from wherever he is to that location. A simple 'Forget' option allows locations to be removed from the list.

There may also be particular applications or services to which the user wishes to have instantaneous access, and/or which may be used at a variety of locations. Such items can be stored in the Vehicle itself and so are always with the user wherever he may be in the world-at-large. The Glove Compartment is located to the side of the navigation controls. When closed, the 'Open' option is displayed. Selecting this causes a moveable window to appear, displaying the contents of the Glove Compartment. Applications are stored in the Glove Compartment by the user dragging their icon from a navigation window onto either the open window or the Glove

Compartment feature on the dashboard. Items are removed by dragging out of the open window and dropping anywhere else.

Views become more interesting when applied in the social sphere. I may want to see only items visited by members of my research team recently. Or I might want to compare one View I have (or my agent has) compiled of interesting sites, with the View a colleague (or his agent) has collected. My View is a way of looking at cyberspace where only things of interest to me exist, and the same applies to him and his View. We can combine these two into another View that shows only those items that are of interest to both of us, or we can create a difference View which shows only those things chosen by only one of us. So the collection of public places that currently comprises cyberspace is filtered to give a socially-shareable and customisable View of cyberspace. This is arguably quite close to the way different groups and individuals hold different views of cities in the real world. The obvious next step is to include representations of cyberspace inhabitants in selective Views. I might want a View that conveys the number of people present in the regions I explore, but I am unlikely to want to see all available information on all the people there. I might want only to see people if they are known to me. I might want to see them differently if they are business colleagues rather than competitors, In general, I will want different attributes of people represented in cyberspace according to their relationship to me. Increasingly, interacting on the Web will become like participating in an on-line multi-user game. Of course, the privacy issues raised are quite daunting.

There are several unanswered questions arising from this work. Is a single hierarchical structure realistic? What are the advantages of Information Islands versus other world models? Would forests, trees, and leaves have been any different (e.g. the “Dataforest”; [14]). Would more than two Views give additional benefits? If we assume a hierarchy of ten Archipelagos, with twenty Islands per Archipelago, twenty Buildings per Island and twenty Floors per Building, we have the necessary scope for a large number of individual information items to be located in the world. With twenty items per Floor, we have 1.6 million items. Relaxing the restrictions on items per Floor by having sub-sets of items accessed by two submenus after the initial Floor Directory selection, and expanding the world to a maximum of 20 Archipelagos, would allow us to accommodate over a billion individual items. Can users navigate in such a world? Almost certainly not if by "navigation" we mean that users can easily find what they want by self-directed wayfinding. The problem of classification is not solved by using 3D, but the nature of the environment is radically changed. Social and personal habitation of large virtual information

spaces may require assessment by criteria other than those we usually apply to information retrieval systems. There is also a good case to be made for the use of an "event horizon" to reduce the complexity of the environment. Waterworth [2] suggests a flat earth metaphor for personal spaces where older items gradually move towards the "edge of the earth", and then disappear.

But in any case this view of expansion is unrealistic. The Information Islands model was designed to meet a particular need. It was assumed that the world would start life relatively empty and would then gradually expand, as providers offered information and other services. In this sense, it is rather like a plan for a city. But it is not clear to what extent development will match the original planning. As Alexander [15] has pointed out, “A city is not a tree” - not a simple hierarchy that grows according to predictable rules. As providers offer services, and users gravitate to the things they are interested in (i.e. willing to spend time and/or money on), the original plans are likely to be heavily modified by market pressures. Like a pleasant city, the world-at-large should evolve to meet the needs of users and providers alike. But not all cities are pleasant, and the balance between central planning and market-led evolution is not easy to strike. The vital question is whether users can find their way around such an evolved model, by whatever means (both agent-mediated and self-directed), to a degree that suits them. Success might be better assessed by the nature of "traveller's tales" [2] reflecting experience of the system, than from more familiar objective measures such as time to "solve a problem" or number of "bad solutions".

4 The Experiential Approach to HCI Design

The problem of interface design has traditionally been characterised as one of communication between the designer and the users. Norman's [16] well-known account of HCI design centres on three kinds of model: the design model (in the head of the designer), the user's model (in the head of the user) and the system image (as presented in the designed interface). The system image serves as the medium of communication between the designer and the user. In the ideal case, the user's model comes to match the design model closely. The common approach to facilitating this process has been to incorporate one or more metaphors in the system image. It then becomes of great importance that the designer chooses appropriate metaphors which convey relevant aspects of the functionality of the system in terms that are understandable to the user [17]. A good metaphor is supposed to permit the user to apply knowledge of the source domain of the metaphor to the unfamiliar target domain of the interface [18].

According to the traditional, objectivist approach to interface design, an interface metaphor is some kind of specialised device for conveying a complex of concepts, based on speaking of, or presenting, one thing as if it were another. However, there is considerable confusion between metaphors and models. A metaphor is not a model, and metaphors are not unambiguous.

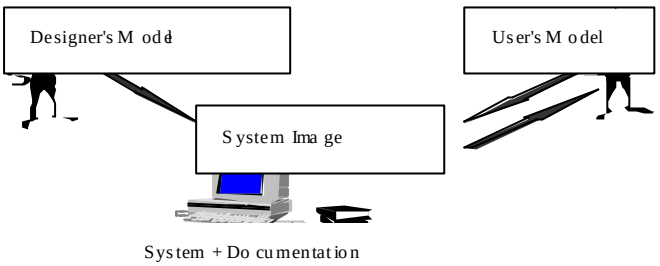

Figure 6 - The traditional view of HCI design

In several books published over the last two decades, George Lakoff and Mark Johnson have presented an alternative view of meaning, one that casts a completely different light on the role and importance of metaphor from that assumed in traditional HCI design [19], [20], [21], [22]. At the same time, they manage to avoid the problems of both objectivism and pure subjectivism.

According to Lakoff and Johnson, metaphor is much more than a linguistic and rhetorical device. They argue that we always think metaphorically, that our everyday experiences are shaped by three kinds of metaphor: structural, orientational and ontological. Structural metaphors are found when one concept is structured in terms of another, for example that argument is war: "Your claims are indefensible", "He attacked every weak point in my argument", "His criticism was right on target", "He shot down all of my arguments", etc.. We not only speak of argument as if it were war (and very pervasively, so that many statements about argument reflect this underlying structuring - although we don't actually think of them as metaphorical), we think about argument as if it were war, and often act according to the same, unconscious, assumption. Orientational metaphors structure experience in terms of spatial orientation. For example, down is negative, up is positive: "I am depressed", "I am really down", "I feel low", "Things are looking up", etc.. Ontological metaphors structure our experiences of abstract phenomena in terms of concrete objects and forces (see also [23]). Efforts to visualise information as shapes, colours, and textures can be seen as reflecting the operation of underlying ontological metaphors, as, arguably, can the application of general interface metaphors such as direct manipulation. The term "synaesthetic media" [24] refers to the presentation of the same information in different modalities, which could be guided by the adoption of appropriate ontological metaphors.

If every concept is metaphorically structured, are we not stuck in some kind of infinite regress? If every concept is structured in terms of another, we are indeed (which is, of course, a fundamental problem with objectivist accounts of meaning). Lakoff and Johnson avoid this infinite regression by suggesting that, at bottom, meaning is rooted in basic, bodily, experiences of life as animals with a certain physical configuration residing on a planet with certain characteristics (notably, gravity). When we use expressions like, "I fell asleep" or "Wake up!" we use metaphor in a way that reflects the physical nature of life on earth. Our body configuration, combined with gravity, makes it necessary for us to sleep in a horizontal position. Johnson [19] provides more detail on the grounding of the (fundamentally metaphorical) conceptual system in corporeal, earthly existence. He proposes the existence of image schemata, which are basic structures of experience. These structures are then projected metaphorically onto more complex experiences. Lakoff ([21], pages 271-278) suggests that image schemata i) are based on bodily experience, (ii) have structural elements, (iii) have a basic logic and (iv) are manifested in actions and expressions. He gives many examples of image schemata, including the container, the centre-periphery, and the verticality schema.

4.1 Designing Experiential Information Landscapes

Our current approach to information landscape design [25], based on this experientialist account of meaning rather than the usual objectivist cognitivism of the traditional "mental model" approach, rests on the fundamental premise that *to design HCI is to design the conditions for possible users' experiences*. In the traditional approach, the metaphor is part of the interface. This need not be the case with experientialism since, by this account, metaphor is everywhere. Taking an experientialist view of interface design suggests that a meaningful interface is one that is experienced in a way that supports the metaphoric projection of image schemata. This is done by the user in the same way that he makes sense of all the other experiences of his daily life, by unconscious projection of bodily image schemata. If the experientialist designer is primarily a creator of user experiences, the traditional interface designer is primarily a communicator of mental models, using metaphor as a useful device.

While I am not arguing that all traditional interface metaphors should be replaced, I do suggest that for several application areas - and these are areas that are at the forefront of current HCI research and development - an experiential approach to HCI design may be more appropriate. A notable example is that of information visualisation and exploration. If we revisit the Information Islands interface wearing our experiential sunglasses, we see that what matters is not so much the

metaphor itself, as the experiential features we chose to take from the real world and incorporate in the virtual.

Conklin [26] argues that "there is no natural topology for an information space", and this claim could be extended to include other aspects of the interactive experience, such as how an information space sounds or how quickly one travels through it. However, an experientialist designer would argue the opposite; that there are, in fact, not one but many natural topologies for such a space, topologies ultimately grounded in human bodily experiences, and projected as image schemata. As mentioned in section 2, Waterworth [2] outlines a design for a Web browsing environment - a Personal Space (StackSpace) - that was informally based around considerations of human bodily experiences in real, physical spaces.

In adopting the experiential approach, a valuable source of design insights is that of language. How do users talk about their experiences? Utterances can be gathered at two stages of the design process: user requirements analysis early on and, later, as corroboration that a particular design is producing the kind of experiences the designer intended. It could be argued that we cannot effectively describe experiences with words but, as Samuel Beckett remarked, they are all we have. The approach to understanding these words is somewhat akin to psychoanalysis; we are looking for the unconscious structures (image schemata) that lie behind the chosen way of describing an experience.

The traditional approach to HCI design uses metaphor to communicate the functionality of the system to the user. The designer draws on users' experiences in another domain to assist their understanding of the system. As Erickson [17] has pointed out, this implies that designers know what the system really is. Despite its problems this approach has been successful in encouraging the widespread use of computers, at least for certain classes of application. The experientialist approach to design also draws on users' prior experiences, but there are several fundamental differences. Firstly, from the traditional perspective, metaphors are useful (usually) but not essential. A traditional user interface metaphor can always be paraphrased into a literal interface. From the experientialist perspective, however, metaphoric projection is essential to the way people make sense of the world, including a user interface. Secondly, that metaphoric projection is essential to sense-making does not mean that we live in a world of metaphors. If we design from an experiential perspective, this does not mean that the interface need be a virtual world of metaphoric objects. Such a world is more likely to be the outcome of the traditional approach.

Experientialism can, however, provide the basic elements of a natural and flexible HCI design pattern language (cf. [27]).

Even though I consider the experientialist approach to user interface design to be a new approach, the experientialist theory of meaning has already attracted attention in fields related to user interface design. For instance, Clay and Wilhelms [28] present a linguistic interface for placement of 3D objects which focuses heavily on spatial relationships as discussed by Lakoff and Johnson. Maglio and Matlock (this volume) demonstrate the usefulness of an experiential analysis in understanding how people conceptualise their explorations of the Web.

There has also been criticism directed towards the experientialist view of metaphor. Although he recognises some merits of experientialism, Coyne [28] claims that Lakoff and Johnson put too strong an emphasis on the primacy of bodily experience and that there are non-embodied and non-spatial uses of concepts like containment and balance. However, Coyne's criticism seems to illustrate, rather than contradict, Lakoff and Johnson's main point; that is, that we project our spatial experiences (embodied as image schemata) to abstract, non-spatial domains of experience.

4.2 SchemaSpace: an experientialist environment

Andreas Lund is currently engaged in a more thoroughgoing attempt at the experiential design of an environment called SchemaSpace [25]. The approach can also be seen as a development of the idea that HCI design is mostly a matter of sensual or perceptual ergonomics rather than the 'cognitive ergonomics' that follow from the traditional, cognitivist approach [23], [24]. SchemaSpace is a three-dimensional virtual environment in which a potential user may organise and browse a collection of references to information sources, located on the Web or elsewhere. As such SchemaSpace is a personal information space [2].

How should an information space like this be designed? We claim that an answer to this question, from a traditional point of view, would in part be formulated in terms of functionality and ways to convey that functionality to the user through the system image. If we instead try to answer the question from an experientialist point of view we first have to reformulate the question as: *what kind of experiences does the user want to get from the interface?* By posing the question this way we put emphasis both on the designer’s role as a creator of meaningful experiences and on the role of the user interface as a source of meaningful experiences.

The intention with SchemaSpace has been to design the interface in such a way that it allows the user to have four different kinds of experiences that each informs the user about different qualitative aspects of the information space:

- *Distinctiveness* - which of the information references belong together, e.g., fall under the same subject or category?
- *Quantity* - how does the number of references in a sub-collection compare to other sub-collections found in the information space?
- *Relevance* - given that a collection of information references belong together, of what relevance is each individual reference in relation to the subject or category?
- *Connectedness* - how do different sub-collections of references relate to each other?

Obviously, these qualitative aspects are by no means all-encompassing and we can see a whole range of other aspects that a user might want to experience from a personal information space. One important dimension not yet addressed by SchemaSpace is that of time; items should show their age, as they tend to do in the real world [2]. However, our purpose here is not to design the ultimate application, but rather to illustrate what we understand to be features of practical experientialist design.

As already mentioned, a meaningful experience is an experience which allows for structuring by means of metaphoric projection of image schemata. Thus, one important step in the design process is to identify image schemata that are associated with the qualitative aspects of the information space we want the user to experience. This identification is by no means arbitrary, on the contrary it ought to be informed by empirical enquiries.

Distinctiveness through containment

In the particular instance of SchemaSpace described here we have about three hundred different Web-references to information on very disparate subjects, ranging from modern literature, via architecture, to computer graphics. Even such a relatively small collection calls for some kind of categorisation, a way to organise and order the information in sub-collections consisting of references belonging to the same category. Put differently, we have to provide for the possibilities of experiencing *distinctiveness*, that is, an experience which informs the user that some information references are in some respect different from other references. In order to provide such an experience we have to identify an image schema which is involved in our general understanding of ordering objects and activities in our everyday life.

Our encounter with containment and boundedness is one of the most pervasive features of our bodily experience. We are intimately aware of our bodies as three-dimensional containers into which we put certain things and out of which other things emerge. Not only are we containers ourselves, but our everyday activities in general - and ordering activities specifically - often involve containment in some respect: we live in containers (houses, shelters, etc.), we organise objects by putting them in different containers. Our frequent bodily experiences of physical boundedness constitute an experiential basis for a *container schema*.

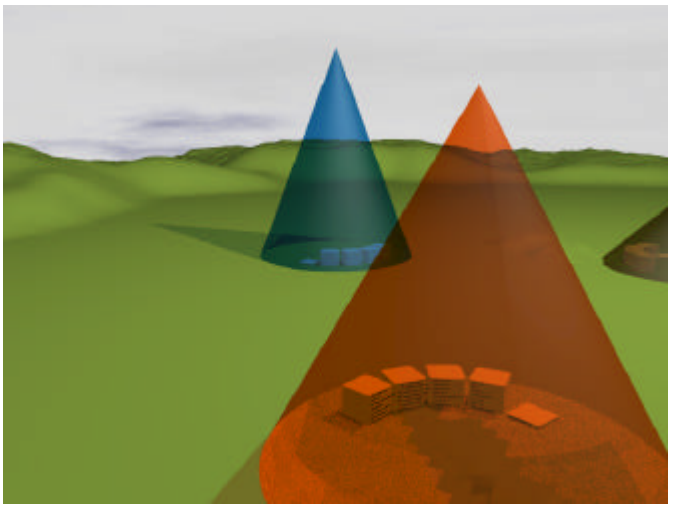

Figure 7 - Distinctiveness through containment and quantity through verticality

A plausible way of providing for the experience of distinctiveness is to present the information references that belong together in a way that allows for a projection of a container schema. There are countless ways of visualising containment and folders and rooms are probably the most familiar user interface containers. However, in our design of SchemaSpace we have as much as possible avoided elements which are - like folders and rooms - heavily metaphorically laden, in order to stress the experientialist features of SchemaSpace (although it is our strong belief that experientialist design need not by necessity exclude ”ordinary” user interface metaphors). Instead, the elements of SchemaSpace consist largely of simple geometric shapes which are not closely associated with a specific source domain. We have chosen to visualise containment by means of semi-transparent cones (see Figure 7). A cone contains information references visualised by stacks of slices (similar to StackSpace; [2]), each with a descriptive textual label. By using semi-transparency it is possible to see that a cone actually contains information references, at the same time as it is apparent that they are bounded by the cone and are thus distinct from other references.

Quantity through verticality

Each cone contains a sub-collection of the totality of information references in SchemaSpace. Some of the sub-collections will contain more or fewer references in comparison to other sub-collections. Even though the cones are semi-transparent, viewed from a distance in the three-dimensional environment it will be difficult to judge the quantity of each cone. In order to provide for a meaningful experience of the quantity of each cone's contents we have to identify an image schema which is involved in our general understanding of quantity. Our basic experiences of quantity are closely associated with verticality (examples from [19]):

> "Whenever we add *more* of a substance - say, water to a glass - the level goes *up*. When we add more objects to a pile, the level *rises*. Remove objects from the pile or water from the glass, and the level goes *down*."

Spatial experiences of the this kind constitute an experiential basis for a *verticality schema*, a schema which by means of metaphoric projection plays an important role in our understanding of non-spatial quantity. Our tendency to conceptualise quantity in terms of verticality reveals itself in everyday language used to talk about quantity:

> "The crime rate kept *rising*. The number of books published *goes up and up* each year. The stock has *fallen* again. You'll get a *higher* interest rate with them. [...]".

In our design we have tried to exploit this verticality aspect of quantity. As shown in Figure 7 the cones in SchemaSpace vary in height. The larger cones have a larger number of references inside compared to the shorter cones. Our intention has been to combine the container and the verticality schema in order not only to express quantity, but also to strengthen the experience of cones as containers of information references.

Degree of relevance through centrality

As already mentioned, one of our goals has been to provide for the experience of distinctiveness. Even if a sub-collection constitutes a unity by virtue of belonging to the same category or subject, different references within a sub-collection may be of different importance or relevance in relation to that particular subject.

As pointed out by Johnson [19]:

> "our world radiates out from *our bodies* as perceptual centers from which we see, hear, touch, taste and smell our world".

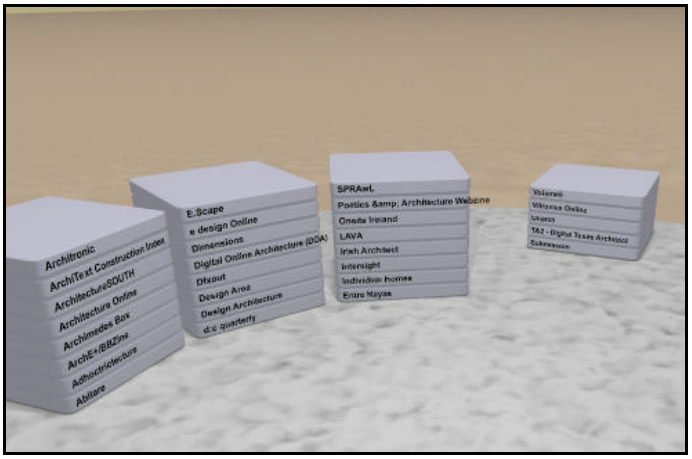

Figure 8 - Degree of relevance through centrality

We also have very basic spatial and physical experiences of centrality as a measure of importance and relevance. Not only is that which is near the centre (the body) within our perceptual reach, but we also experience our bodies as having a centre and periphery where the central parts (trunk, heart, etc.) are of greatest importance to our well-being and identity [21].

In order for a potential user of SchemaSpace to experience some references as more important and relevant in relation to other references within a cone, we exploit a *centre-periphery schema*, which has its experiential grounding in perceptual experiences of centrality mentioned above.

As seen in Figure 8, stacks of information references are organised along an arc. In those cases where there are a lot of references within a cone, the arc will eventually be closed and form a circle centred around the vertical axis of the cone. Information references can, however, be placed at varying distances from the centre; that is, some references will perceptually be closer to the centre and some will be more peripheral (see stack to the right in Figure 8). Our goal with this arrangement is to invoke a metaphoric projection - on the part of the user - of the centre-periphery schema in order to experience those references which are perceptually central as conceptually central.

Connectedness through linkage

Finally, we want the user to experience that some sub-collections of references are related to each other, even though they are distinct from each other. In SchemaSpace we have a collection of references on the subject Virtual Reality Modelling Language (VRML). But we also have two categories with references to information on VRML-browsers and VRML-worlds. These two categories may be considered as distinct from VRML information in general, but not in the same sense as information on the writer Paul Auster is distinct from information on architectural

magazines. There is a connection between general information on VRML and VRML-worlds and browsers, that does not exist in any obvious way between Paul Auster and architecture. In order to provide for an experience of this kind of connectedness we exploit a *link schema*. The link schema is often involved in our understanding of relations and connections of different kinds, not only physical connections, but also more abstract, non-physical connections like interpersonal relationships.

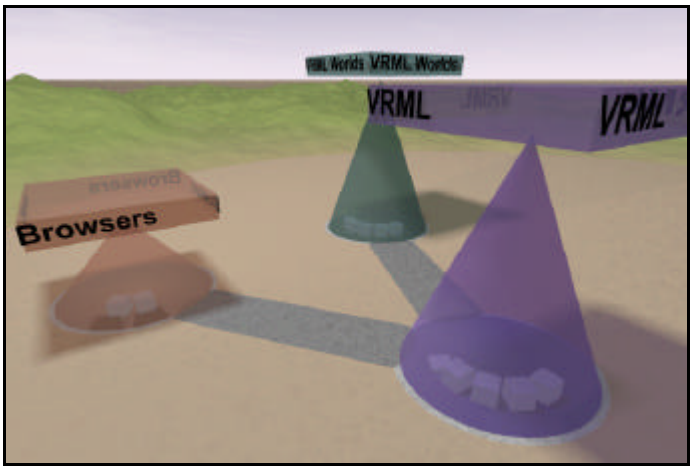

Figure 9 - Connectedness through linkage

In SchemaSpace cones are connected with a path-like link if the sub-collections contained in the cone are considered to be connected, as is the case with VRML in general and VRML-browsers and worlds (see Figure 9). As with the centre-periphery example above, our goal with this arrangement is to provide the user with perceptual cues that allow for structuring by means of projection of a certain schema, in this case the link schema. Of course, linking to show connectedness depends on some knowledge of what is connected and why. This can be approached in several ways which will not be described here.

From the experientialist view, what is needed in HCI design is for the interface to be a source of experiences, designed in such a way that the experiences generated may be structured by the projection of basic bodily image schemata. As in poetry, metaphor is used to create an effect in the experiencer. What the resultant interface (or poem) means, *what it is* for a given user, depends on his or her unconscious reactions to the structures provided. If the interface feels right for its purpose, it is successful. No designer can know what the system really is, in general. It is what it means to individual users (often as members of social groups) and, like life, it means what it is experienced to be.

5 Presences and Concealment

Even if we don't want to be personally identifiable on the Web (in the same way we sometimes don't want to be individually known to strangers in the real world) we currently don't even have presence as anonymous people in cyberspace. All we can tell, sometimes, is how many people have visited a site before us. Or rather, how many visits have been made to the site. For cyberspace to become real, we need a sense of people's presence (and absence), with suitable protection for privacy - if that is possible.

A limited sense of what shared presence in the Web would bring is provided by experiences in "multi-user dungeons" (MUDs) and the Internet Relay Chat services. The 1995 book by Turkle [30] gives an insight into those worlds, although the MUD and Chat users are probably not typical of Web users. Specifically, they are self-selected for their interest in role playing and/or a need to alleviate real life loneliness.

However, these environments, insofar as they use text descriptions of space rather than 3D graphics, are also quite unlike a shared, as-if-real, virtual world. As already mentioned, a textually-described world of the kind investigated by Dieberger (e.g. this volume) is an unlikely candidate as a medium for shared exploration of information space. This is because conceptualising and navigating such a space is so demanding that little mental capacity remains actually to deal with the information located during navigation. The trend in the navigation of information spaces, and indeed in HCI in general, is to shift the burden of dealing with the environment from conceptual, linguistic processing and conscious decision making to direct perceptual processing within automatised sensori-motor behaviours. In other words, to allow people to deal with information while minimising the extent to which they have to deal consciously with information about how to deal with information.

Turkle [30] points to the ease of adopting multiple personae in cyberspace, to present the face we choose to present rather than the real-life person we have become over the years. This can be seen as partial or selective presence. We can think of degrees of presence, from totally concealed (invisible), through anonymous (featureless) but visible, to articulated personae one of which might be a representation of our real world personality.

Should we be able to choose how we appear to others? Should we be able to appear present when we are not, and not present when we are? False presence arises when we appear to be somewhere, but are actually elsewhere. Each of us wants to know as much as possible about others, but to control what is known about ourselves. We achieve this in the real world by limiting

the time we spend, or the visibility we have, in public places, retreating to personal spaces when we feel the need. Multiple personae multiply the scope for deception, and the creation of personae is, of course, much easier in the virtual world than the real one.

As more and more people migrate to cyberspace, both the amount of information and the number of sources of information multiply. But human attention is still singular. We are each aware of only one thing at once. I can watch TV or read a book, but I cannot do both at any particular moment (of course, I can switch between the two, and if the rate of information transfer from the TV is typically low I won't miss much).

In the same way, even though many browsers now offer some support for following multiple threads, we can only really attend to one thing at once, even with multiple display windows. This is how we can be said to follow links - to navigate - at all, and to be lost in cyberspace when we lose our (singular) way. This also helps give us the presence that is currently unrepresented in the Web. In cyberspace we are where our attention is focused, but we have no presence until we are visible in public.

6 Conclusions

The Web differs in several major respects from the hypermedia systems that were the focus of so much premature research in the 1980s. Like a capital city in a developing country, it is large and growing very rapidly, both in the amount of information available and the number of inhabitants. All those people use the same system, rather than having their own copy; in other words, they truly co-habit cyberspace. And no-one is designing it as a whole. Rather, we operate "locally" by introducing innovations that may or may not catch on in the electronic world-at-large. The design of such landscapes and features is more appropriately based on notions of meaning as experience, rather than traditional ideas of meaning as functionality conveyed through HCI models, since we do not know what the function of shared information landscapes might be. But such a situation is natural for us, since we don't know what the function of the real world is.

Individuals need personal space to make sense of the information they collect from the world at large. People (like all animals) are naturally equipped to deal with 3D spatial environments without that imposing a heavy burden on their scant attentional resources. In fact, 3D interactive environments are such a powerful way of presenting collections of information just because they allow people to explore virtual worlds of information in the same ways they explore the real world, whether this is an individual activity or in groups. Views of spaces can serve as a powerful

mechanism for social interaction because they can be compared, contrasted, and exchanged. Groups interacting in real time require a sense of presence in shared 3D space.

In the future, people will be represented with degrees of presence, and we will search for each other in cyberspace, not only for the things we have created. This will allow us to behave more naturally as the social animals we usually are. Personae will multiply but attention (and thus true presence) is still singular. To be truly present implies communicable with, and visible. When we have such presence the Web will gain social context naturally.

Experiential design captures basic, unconscious, animal reactions to physical environments and introduces them to shared virtual landscapes. We can design appropriate tools and environments, just as we can design churches, cinemas and houses, but we do not design societies or social behaviour. We are social (and spatial) by nature, not design.

7 Acknowledgements

Andreas Lund has done all the work on SchemaSpace, and has made large contributions to the sections of this chapter that deal with the experiential approach to HCI design. Eva Lindh provided insightful comments that helped me make the text less unclear. David Modjeska, on secondment from the University of Toronto, has provided useful feedback and stimulating discussions on this text and on related research issues.

8 References

1. Waterworth, J. A. and Chignell, M. H. (1989). A Manifesto for Hypermedia Usability Research. *Hypermedia, 1,* (3), 205-2

2. Waterworth, J. A. (1997a). Personal Spaces: 3D Spatial Worlds for Information Exploration, Organisation and Communication. In R. Earnshaw and J. Vince (eds.): *The Internet in 3D: Information, Images, and Interaction*. New York: Academic Press, 1997.

3. Harrison, S. and Dourish, P. (1996). Re-Place-ing Space: The Roles of Place and Space in Collaborative Systems. Proceedings of ACM Conference on CSCW'96, Boston, November 1996.

4. Meyrowitz, J. (1985). *No Sense of Place.* New York: Oxford University Press.

5. Waterworth, J. A. (1992). *Multimedia Interaction: Human Factors Aspects.* Chichester, UK: Ellis Horwood (Simon and Schuster, 199.2

6. Dieberger, A. (1993). The Information City - a step towards merging of hypertext and virtual reality. Poster at *Hypertext '93*.

7. Waterworth, J. A. and Singh, G. (1994). Information Islands: Private Views of Public Places. In *Proceedings of MHVR'94 East-West International Conference on Multimedia, Hypermedia and Virtual Reality.* Moscow, September 14-16, 1994.

8. Waterworth, J. A. (1995). Viewing Others and Others' Views: Presence and Concealment in Shared Hyperspace. Presented at *Workshop on Social Contexts of Hypermedia*, 16-17 February 1995, Department of Informatics, Umeå University, Sweden.

9. Waterworth, J. A. (1996a). A Pattern of Islands: Exploring Public Information Space in a Private Vehicle. In Brusilovsky, P, Kommers, P and Streitz, N (eds.) *Multimedia, Hypermedia and Virtual Reality.* Springer Verlag Lecture Notes in Computer Science, 1996.

10. Musil, S and Pigel G (1993). Virgets: Elements for Building 3-D User Interfaces. In *Proceedings of the Symposium Virtual Reality Vienna*, December 1-3, 1993. Also available as TR 93/13, Vienna User Interface Group, Lenaugasse 2/8, A-1080 Vienna.

11. Henderson, D A and Card, S K (1986). Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention in a Window-Based Graphical User Interface. *ACM Transactions on Graphics*, 5 (3), 211-243.

12. Nagel, T (1986). *The View from Nowhere*. New York: Oxford University Press.

13. Fishkin, K. and Stone, M. C. (1995). Enhanced Dynamic Queries via Movable Filters. *Proceedings of CHI'95*. New York: ACM.

14. Rifas, L (1994). The Dataforest: tree forms as information display graphics. In Dieberger, 1994, Report of the Workshop on Spatial Metaphors at *ECHT'94 - the European Conference on Hypermedia Technology*, September 1994, Edinburgh, UK..

15. Alexander, C. (1982). A City is not a Tree. In Kaplan, S. and Kaplan, R. (eds.) *Humanscape - Environments for People.* Ann Arbor, USA: Ulrich's Books, pages 377-402.

16. Norman, D. (1986). Cognitive Engineering. In Norman, D and Draper, S., *User Centered System Design*. Hillsdale: Lawrence Erlbaum Associates.

17. Erickson, T. D. (1990). Working with Interface Metaphors. In Laurel, B (ed.), *The Art of HCI Design*. Menlo Park, USA: Addison-Wesley.

18. Gentner, D., Falkenhainer, B. and Skorstad, J. (1988). Viewing metaphor as analogy. In *Analogical Reasoning: Perspectives of Artificial Intelligence, Cognitive Science and Philosophy*. D. H. Helma (ed.), Kluwer.

19. Johnson, M. (1987). *The Body in the Mind.* Chicago: Chicago University Press.

20. Johnson, M. (1993). *Moral Imagination.* Chicago: Chicago University Press.

21. Lakoff, G. (1987). *Women, Fire and Dangerous Things.* Chicago: Chicago Univ. Press.

22. Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By.* Chicago: Chicago University Press.

23. Waterworth, J. A. (1996b*).* Virtual Reality for Animals. *Proceedings of Ciber@RT'96,* First International Conference on Virtual Reality. Valencia, Spain, November 1996.

24. Waterworth, J. A. (1997b). Creativity and Sensation: The Case for Synaesthetic Media. *Leonardo, 30,* (3).

25. Lund, A. and Waterworth, J. A. (1998). Experiential Design: Reflecting Embodiment at the Interface. *Computation for Metaphors, Analogy and Agents: An International Workshop*, University of Aizu, Japan, April 1998.

26. Conklin, J.(1987). Hypertext: An Introduction and Survey. *IEEE Computer, 20* (9), 17-41.

27. Alexander, C., Ishikawa, S., Silverstein, M. (1977*). A Pattern Language: Towns, Buildings, Construction.* New York: Oxford University Press.

28. Clay, S R. and Wilhelms, J. (1996). Put: Language-Based Interactive Manipulation of Objects. In *IEEE Computer Graphics and Applications*, Vol. 16, No. 3, May 1996.

29. Coyne, R. (1995). *Designing information technology in the postmodern age: From method to metaphor*. Cambridge MA: MIT Press.

30. Turkle, S. (1995). *Life on the Screen*. New York: Simon and Schuster