Chapter Two:
Wayfinding in the Real and Virtual World

Abstract

This chapter will focus upon the body of literature concerned with wayfinding in the real world and will review recent research undertaken into wayfinding in the virtual world. It will begin by discussing a selection of definitions of wayfinding including an examination of the origins of the word. Through this survey of definitions, this chapter will gradually distil a working definition of wayfinding to be used throughout this thesis. Next, in reviewing wayfinding in the real world, different methods of assessing wayfinding performance are presented and assessed. Wayfinding papers that focus upon the effect of the environment on wayfinding performance are then reviewed and their methods discussed. After identifying problems with both these aspects of wayfinding research (the assessment of behaviour and effect of environment), this chapter goes on to examine more recent work. This work investigates wayfinding performance in the virtual world, and asks whether virtual environments are considered adequate research tools to investigate this phenomenon. Papers highlighting issues of research methodology are reviewed, with particular attention paid to papers seeking to establish whether interface and/or procedure have any effect on resultant data. A series of wayfinding experiments conducted solely in a virtual environment are then discussed, highlighting factors that lead to a series of papers that investigate the effect of Lynch-inspired environmental components (landmarks, paths, edges, nodes and districts) upon wayfinding. Finally, a number of papers attempting to compare real and virtual wayfinding behaviour are compared, leading to the conclusion that, broadly speaking, the same approach is taken by all the researchers reviewed. Rather than wayfinding behaviour, it is resultant spatial knowledge that is being analysed. Assumptions of equivalence (that real wayfinding correlates to virtual wayfinding) are made based solely on this. This chapter concludes with an observation that more objective methods of measuring wayfinding performance are needed, coupled with better analyses of environments, and that the degree to which real and virtual wayfinding performance are analogous still needs to be established.

Introduction

The focus of this Chapter will be the examination of the body of knowledge of wayfinding, both in the real world and, from recent years, in the virtual realm. If the research question underpinning this thesis is whether it is possible to *learn* from the study of virtual environments *how* people will behave in real environments, then it is vital to first understand what is *already* known about behaviour in real environments. In particular, the type of real-world behaviour of greatest relevance to this thesis is the act of what has come to be termed 'wayfinding'.

What is Wayfinding?

Arthur and Passini, in (Arthur and Passini 1992) attribute the term 'way-finding[1]' to Lynch, stating that its first occurrence was used in the book, The Image of the City (Lynch 1960). However, they estimate that it did not come into widespread use until the late 1970s, when it essentially replaced the phrase 'spatial orientation', hitherto used in academic writing. They attempted to seek some derivation and hence justification for the term wayfinding by examining its etymology, assuming it to be a derivation of the words wayfarer and wayfaring. Both of these words are derived from Old English[2]; wayfaring (archaic) was first recorded as being used in 1536 AD, whereas an older version of the word, wayfering (obsolete) can be traced back to 890 AD[3]. The definitions of these words mean journeying or travelling, particularly on foot.

Another term that may have some bearing on the phrase wayfinding is the term pathfinder. Pathfinder is a word of North American origins (although similar in sound and meaning to the modern German word, pfadfinder), whose usage can be traced back to the middle of the Nineteenth Century, indicating a person who discovers a path or way; an explorer. The origin of this word has been attributed to the author James Fenimore Cooper, as it formed the title of his book, “The Pathfinder”, published in 1840, a term that was subsequently taken up by the English[4]. If the words pathfinder and wayfarer were to be combined, then the resultant term *wayfinder* could be held to be an obvious hybrid of the two parent words (its derivations partly Old English, partly Modern American). If one were to attempt to blend the *definitions* of the two words in a similar manner, one could readily surmise that the term wayfinding should describe a process of *travelling on foot in an exploratory manner*. Lynch's modern coinage of the word has unfortunately yet to find its way into any of the dictionaries of modern usage (either English or American English) even though it has now been in existence for more than forty years and has been used in the title of at least two books (Arthur and Passini 1992; Passini 1992).

If wayfinding is a relatively new word with scarcely any historical precedence, to what does it refer, and have any commonly agreed definitions, been established? At the very simplest end of the spectrum of definitions, as stated in (Carpenter 1989), “Wayfinding refers to what you do to find your way somewhere.” However, this is perhaps a little too oversimplified and gives no indication of what it is that is actually being *done*, in order to find one's way. In (Arthur and Passini 1992) the concept of problem solving is introduced into a definition. They start with three distinct phrases that can be amalgamated into a single description; “Wayfinding is continuous, spatial problem solving under uncertainty.” Although this definition begins to give some indication of what people do when they navigate from an origin to a destination, the act of "continuous, spatial problem solving under uncertainty" is still extremely broad and vague. For example, the act of playing the popular computer game, Tetris[5] would fall neatly into this definition. It is clear that a definition is required that implies not only the act of travelling from origin to destination plus the act of spatial problem solving (potentially including the *types* of problems to be solved) and encompasses a person's cognition of their environment.

If definitions of wayfinding that include the concept of environment-perception or cognition are now examined, it is necessary to initially return to Lynch, whose original definition of wayfinding was "a consistent use and organization of definite sensory cues from the external environment." In this definition Lynch stresses the importance of our senses to the act of wayfinding, yet omits to describe *how* it is that we use this information. In this respect, he is omitting the first two aspects of wayfinding mentioned in previous definition.

However, the primary importance of the input of our senses (our perception and cognition of the environment through which we navigate) is evident from the number of definitions of wayfinding which concentrate upon this aspect. For example, Gibson, in (Gibson 1979) stresses that " purposive locomotion such as homing, migrating, finding one's way [wayfinding], getting from place to place, and being orientated, depends on just the kind of sequential optical information [continuous visual perception of the environment] described." Although Gibson never states what wayfinding is, he is quite clear about what it is not. With regard to the theory of response chains and cognitive maps he says, "Wayfinding is surely not a sequence of turning responses conditioned to stimuli. But neither is it the consulting of an internal map of the maze" Although Gibson stresses the importance of perception and specifically visual perception to the act of wayfinding, he also concedes that there is more to wayfinding than purely responding to visual information in the environment. Namely, that although cognition and perception are essential to a definition of wayfinding, on their own they are insufficient.

If we wish to find an alternative example of a definition of wayfinding, which stems from a focus upon *only* the visual aspects of wayfinding, we come to a definition by Cutting, from (Cutting 1996), where he states that wayfinding is "how people find their way through cluttered environments with ease and without injury." This definition makes no reference to the acts of exploration or route finding. This is less than surprising since Cutting's research is into retinal optic flow. It becomes evident therefore, that the particular emphasis of various definitions can differ depending on the academic discipline of the author. For this reason it is suggested that an attempt to find or coin a single definition of wayfinding that is acceptable to a range of academic disciplines may be an impossible task.

In a manner similar to Gibson and Cutting, Golledge also emphasises the relationship between navigation and vision in his definition in (Golledge 1995), where he states that wayfinding "appears to be one of the primary functions of vision in virtually all biological systems. The processes involved includes cue or landmark recognition, turn angle estimation and reproduction, route link sequencing, network comprehension, frame of reference identification, route plotting strategies (e.g. dead reckoning, path integration, environmental simplification and en-route choice, shortcutting)." In this statement Golledge takes a rather unusual standpoint. He not only states that visual perception is *necessary* for wayfinding or navigation, he actually inverts this causal relationship and suggests that the necessity to move through our environment is one of the principal reasons for our sense of vision (i.e. we do not wayfind because we can see; we see so that we can wayfind.). In his definition he lists some of the processes he believes we perform whilst navigating, and although he does not mention spatial problem solving as a general term, many of the processes he describes can be regarded as spatial problem solving (others are clearly not, such as landmark recognition). In this way he begins to combine the acts of environment cognition and spatial problem solving into a single definition of wayfinding (albeit a task-list definition). This is a useful starting point for establishing what the 'do' of Carpenter's definition ("Wayfinding refers to what you do to find your way somewhere.") may really entail.

To continue this line of enquiry, it is possible to turn once again to Arthur and Passini in (Arthur and Passini 1992), where they state in another definition that wayfinding is "all the perceptual, cognitive, and decision-making processes necessary to find one's way." In this definition, Arthur and Passini, like Golledge above, are introducing cognitive and perceptual aspects of wayfinding into their definition, combined with a recognition of the necessity to make reference to the act of making decisions (spatial problem solving). Arthur and Passini eventually formulate a final definition of wayfinding, which they include in their glossary of wayfinding terms in (Arthur and Passini 1992). This definition states that wayfinding consists of "finding one's way to a destination; spatial problem solving comprising three interdependent processes: decision making, decision executing, and information processing." In this definition, they not only combine the act of travelling from origin to destination under uncertainty with the need for spatial problem solving, but in addition the perceptual and cognitive aspects are partially covered by the phrase "information processing".

Before completing this section it is worth briefly noting a couple of recent definitions of wayfinding that are somewhat at odds with the majority of definitions noted in this Chapter. In (Darken, Allard et al. 1999) and (Bowman 1999) both authors refer to wayfinding as a term that describes solely the *cognitive processes* involved in finding one's way from an origin to a destination. Darken, in the editor's introduction to the above issue of the journal Presence (Darken, Allard et al. 1999), says, "We know that what we often refer to as *navigation* is not merely physical translation through a space, termed *locomotion* or *travel*, but that there is also a cognitive element, often referred to as *wayfinding*, that involves issues such as mental representations, route planning, and distance estimation." Later in the same issue, Bowman makes the same distinction when he says, "We define *navigation* as the complete process of moving through an environment. Navigation has two parts: *wayfinding* (the cognitive decision-making process by which a movement is planned), and *travel* (the actual motion from the current location to the new location)." In both of these definitions, their use of the word navigation is far closer to other academic definitions of wayfinding, especially Passini's. However such definitions appear not to be used elsewhere in other texts on wayfinding. One argument that could be used to counter against such a definition is that cognitive processes, considered in isolation of any movement through the environment, would be fundamentally meaningless, a standpoint first presented by Gibson. Therefore, since the act of travelling through an environment is a prerequisite component of our cognition of that environment, then the act of wayfinding must encompass both movement and cognition.

In order to arrive at a final, useable definition of wayfinding, it is necessary to consider briefly a selection of definitions that include a reference to spatial decision-making. In particular it may be useful to examine the relationship between environmental cognition and decision making. Passini says in (Passini 1992) that the "decision execution process based on a matching feedback mechanism in which expected environmental information… is matched with perceived information." This reference to a feedback mechanism has some similarities to a much earlier definition of wayfinding by Downs and Stea in (Downs and Stea 1973). In this book, they define wayfinding as comprising four stages. The first stage is that of orientation or the determination of both self-location and target-location (or estimated target-location) within the environment. The second stage is initial route choice, the selection of a route from starting location to target location. The third stage actually runs in parallel with a series of recursive instances of stages one and two. Stage three is route monitoring, constant checking of the route taken, modified by estimates of self-location and target-location (stage one) and reassessment or confirmation of route choice (stage two). The final stage of the process is the ability to recognise when the target has actually been reached. This can be summed up by the following diagram, in figure 2.1.

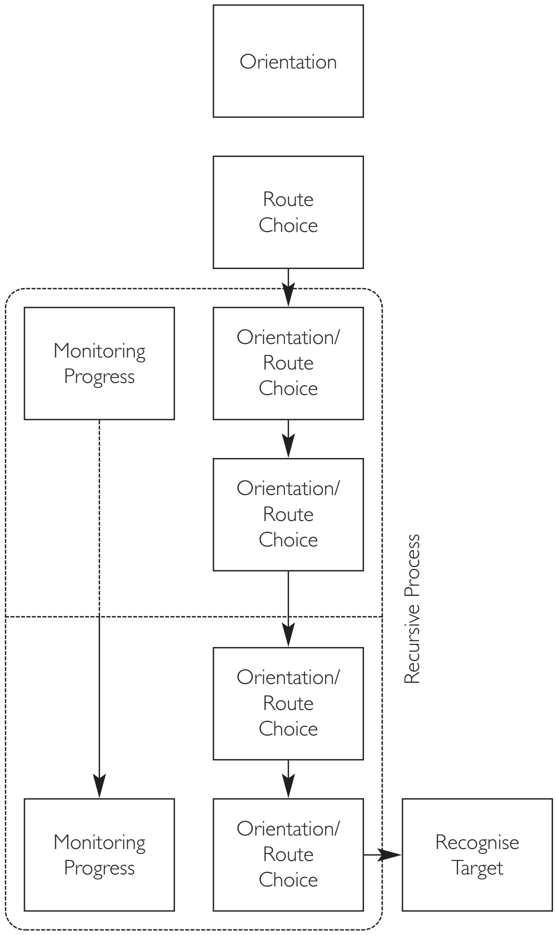

Figure 2.1 The Process of Wayfinding, After Downs and Stea

To summarise, the factor that appears to be essential to Downs and Stea's definition of wayfinding is a recursive act of measuring all decisions taken against continuous environmental cognition (monitoring progress). If this concept of recursive actions is taken into account, it can perhaps be used to coin a new definition of wayfinding, a definition that attempts to combine many of the components of other definitions surveyed in this chapter. One definition that could be thus created is that wayfinding is *the act of travelling to a destination by a continuous, recursive process of making route-choices whilst evaluating previous spatial decisions against constant cognition of the environment.* This is the definition of wayfinding to be used in this thesis. However, faced with such an all-encompassing and complex definition, it is probably worth being reminding ourselves that wayfinding is, in Raubal's words, "a basic activity that people do throughout their entire lives as they navigate from one place to another." (Raubal, Egenhofer et al. 1997).

Wayfinding in the Real World

In their paper (Raubal and Worboys 1999), Raubal and Worboys make another useful observation about wayfinding and the distinction between "knowledge in the head" and "knowledge in the world". These concepts originated from Norman in (Norman 1988). Norman makes these comments about the two kinds of knowledge, "Knowledge (or information) in the world and in the head are both essential in our daily functioning… Knowledge in the world acts as its own reminder. It can help us recover structures that we otherwise would forget. Knowledge in the head is efficient: no search and interpretation of the environment is required… Knowledge in the world is easier to learn, but often more difficult to use, And it relies heavily upon the continued physical presence of the information. Once the environment is changed then the information available is also changed. Performance relies on the physical presence of the task environment." He goes on to say, "Because you know that the information is available in the environment, the information you internally code in memory need be precise enough only to sustain the quality of behaviour you desire. This is one reason people can function well in their environment and still be unable to describe what they do. For example, a person can travel accurately through a city without being able to describe the route precisely."

If we consider these two types of knowledge in the context of the act of wayfinding, it is clear that the knowledge in the world is present in many forms and at many cognitive levels. At a lowest level of awareness, this knowledge can be regarded as being implicit in the overall configuration and structure of the environment, whereas at a higher level of awareness the knowledge in the world is explicit in the forms of, for example, signage. This range of 'knowledge in the world' (the range from implicit to explicit) can also be regarded as an example of non-discursive versus discursive environmental cues. One reason why subjects find it so difficult to describe the visual cues they use when wayfinding is that they only have words to name the explicit examples of knowledge in the world. It is relatively easy to determine whether a subject has attended to and/or acted upon information provided in the form of a sign (for example, we can ask them). It is difficult to identify a visual cue for which you have no concept or no name. As Peponis and Zimring say in (Peponis, Zimring et al. 1990), "the way that people verbalise or draw their understanding of the environment may be quite different from how they actually conceptualize it or navigate through it." Therefore, one of the difficulties inherent in examining the 'knowledge in the world', is that we can readily identify some visual cues, yet find it difficult to identify others. In contrast, in wayfinding terms, 'knowledge in the head' may be regarded as strategy, deliberate actions/decisions, and applications of past experience and memory.

Raubal and Worboys quite rightly level a criticism at wayfinding research, namely that all too often, research has focussed on the 'knowledge in the head' whilst ignoring the 'knowledge in the world'. Another way of phrasing this is that researchers into wayfinding have concentrated their efforts upon aspects of their subject's task performance in order to determine the effect of the world, rather than analysing the environments through which the subjects navigate. Rarely have both of these factors been examined equally. This is an interesting observation, since 'knowledge in the head' is patently the more impenetrable of the two areas of knowledge, as it is by far the more difficult to gauge. We are mostly able to assess 'knowledge in the head' by examining secondary sources of information only - what people *say* they did or are doing. It is almost impossible to examine primary sources of information, namely to examine what is *actually* happening in the brain whilst navigating. The closest attempt to do this is the work done by O'Keefe on hippocampal activity during wayfinding. A tertiary source of information is available by examining people's behaviour to see whether their actions shed any light on their 'knowledge in the head'.

However, despite Raubal and Worboys' criticism, it is possible to level a far more serious criticism at this body of research. Irrespective of whether 'knowledge in the head' or in the 'world' is being investigated, the majority of research in the field has suffered from a paucity of objective research methods. In order for work in this area to progress it is vital that researchers find more objective ways to analyse the 'knowledge in the head', i.e. the thoughts and actions of people in the environment. Moreover, it is necessary to seek more objective ways of analysing the environments themselves, so that we can determine which are the salient qualities of the environment that influence wayfinding performance. In order to begin to suggest ways in which research into the effect of these two phenomena can be furthered, it is necessary to examine research methods that have been used in the past. Some of the following papers focus upon knowledge in the head (both subjectively and occasionally objectively) and some focus upon knowledge in the world (again both subjectively and objectively).

This chapter section will start by examining how researchers in the field have attempted to investigate the role that 'knowledge in the head' plays in wayfinding. One approach taken to reach this goal is to analyse wayfinding performance. At the more subjective level, a common method is to ask the subject whether or not they found a wayfinding task easy or difficult. As Peponis et al. say in (Peponis, Zimring et al. 1990), "Direct observation of wayfinding is relatively rare, and even when it is attempted it is not always clear what is being recorded and how it might be analyzed." For example, although Braaksma and Cook in their paper, (Braaksma and Cook 1980), present an objective method for analysing an environment, their method for comparing this analysis to wayfinding performance is achieved by conducting informal interviews. They ask people to estimate the ease with which they found their way in the environment (in this case an airport).

Passini, in (Passini 1992), also uses post-test interview evaluations. In his case examples, he asked subjects to assess their own performance of the wayfinding tasks, as well as assessing the settings and the signage system (environmental factors). Asking people to gauge whether or not they feel an environment aids or hinders wayfinding, or to assess how well they judge that a wayfinding task was completed, is particularly problematic. These are the most subjective measures of all possible methods of assessing wayfinding performance.

Passini goes on to describe another method he uses to assess the performance of his subjects' wayfinding. This is a method that has subsequently been reproduced by a number of researchers, for example Raubal in (Raubal, Egenhofer et al. 1997), (Raubal and Egenhofer 1998) and (Raubal and Worboys 1999). Passini defines this method as eliciting "a wayfinding protocol", namely the "verbalization of the subject's problem-solving process while completing the wayfinding task." In effect, a researcher follows a subject, whilst that subject performs a wayfinding task. The subject's continuous, verbal commentary is recorded and later transcribed by the researcher. Passini describes the method used to prompt the subjects to verbalise their experiences. "The subjects were encouraged to describe freely what went through their minds at any time, to discuss the decisions they made, and to indicate what information they relied upon. The function of the investigator was to encourage this verbalization during the walk. The protocol contains an unedited version of the whole conversation pertinent to the topic as it was recorded on tape." Of course, the primary problem with this method, a problem that Passini was well aware of, was that the very presence of the interviewer, accompanying the subject through the environment, might have an effect on the performance of the subject.

Passini was particularly concerned about the researcher giving involuntary clues through their body language or facial expressions. He describes this problem thus, "After a short training period it did not appear difficult to follow a subject passively in his choice of walking directions, walking speed and so on, but it did prove difficult to refrain from behavioural reactions prompted by unexpected events."

The other problem with this method is a problem that Passini seems not to have considered, and is linked to the issue (touched upon earlier in this chapter) of identifying *what* knowledge in the world is important. For example, during this process, subjects are not only encouraged to describe what they are doing, and decisions they have made, for example, "That's not the direction I want to go". Subjects are also prompted to describe visual cues that they pick up, such as; "I see a sign that says I should go down that hall to go to gate A[6]". It is a relatively straightforward task to comment that an information-sign has been noted, but less easy to put into words the effect that, for example, the shape of the space being occupied has upon decision-making. Occasionally spatial attributes are reported, for example, "I come out in a big taller area[7]", "it's a long space", "a clear open area" and "The space is much narrower[8]" However, the vocabulary we have at our disposal to describe such phenomena is deficient. We do not have the requisite tools to describe precisely what aspect of the environment is of influence. In addition to this, there exists the additional danger that the subjects are merely putting into words, what they think the interviewer wants to hear, and hence the 'experimenter effect' comes into play. For these reasons, the method of asking subjects to give a running commentary ultimately remains an all too subjective and hence less than satisfactory method of investigation.

Another member of the family of 'self-evaluation' methods is the questionnaire used by Weisman in (Weisman 1981) to judge the extent of wayfinding problems in a set of ten different buildings. In this questionnaire, Weisman poses questions querying the number of times and degrees to which a subject has been 'lost' in the building and whether they would feel confident to give route directions to a stranger. This type of questionnaire is a popular method of self-assessment (and at least more directed than an informal interview). Nevertheless, Weisman, himself, recognises the problems inherent in this approach and at the end of his paper calls for the development of a set of more objective measures of wayfinding behaviour.

Towards the end of Weisman's questionnaire, the subject is asked to draw a diagram of the layout of the building. This method of assessing the knowledge of the building is taken directly from Lynch (Lynch 1960) and is referred to as a cognitive map. Lynch uses cognitive maps to determine the extent of a subject's knowledge of their environment. One problem with this approach was summed up succinctly by Passini and Arthur in (Arthur and Passini 1992), where they comment that, "The search for an answer [to the relevance of cognitive maps] was dampened by the observation that in many situations people got around quite well and did not feel disorientated even if they had only a very rudimentary understanding of the setting." In other words, there is no indication that an ability to draw a diagram of a building or environment necessarily has any bearing on wayfinding ability and vice versa. Finally, an additional problem with this method is how to begin to assess the qualities of a map so drawn. Gärling tries to explain this in (Gärling, Böök et al. 1986) when he says that there are doubts about "whether the results [of drawing cognitive maps] can be generalized" and that "problems of reliable scoring [assessing the results] are nevertheless likely to arise"

The criticisms cited above can also be levelled at another method of assessing wayfinding ability. An established method of assessing wayfinding ability is by using scene recognition tests, as described by Gärling in (Gärling, Böök et al. 1986). Once a wayfinding task has been completed, a subject is presented with a number of pictures or photographs of various scenes and asked to identify which of them had been encountered along their route. Again, a successful ability to perform this type of memory-recall test need not necessarily be indicative of an ability to have performed the original wayfinding task well.

Finally, a method that is also described by Gärling in the same paper is a particularly well known and popular method used extensively in wayfinding research. At various stages throughout a journey, participating subjects are instructed to point in the direction of salient locations or landmarks in the environment and/or to gauge their distance from the subject's location. The errors between the subject's estimations and their true orientation/distance are then calculated. The errors for all subjects in the test are then compared. This particular test has the advantage of being a more *objective* measure than the majority of those described so far. The relative objectivity of this technique has obviously contributed to its popularity as a research technique. However, this test essentially measures little more than a subject's estimation of their own position within an environment plus the relative position of other locations. Gärling comments that "The ability to localize reference points in the environment… [is] an important factor in the ability to maintain orientation." However, an ability to perform this task well is again not necessarily indicative of other abilities. Although clearly a useful measure in its own right, the question that must be asked is the same question that Passini raises with regard to cognitive maps. Namely, is this skill necessarily vital to the ability to find one's way through an environment or is it simply indicative of an innate ability to estimate orientations and distances? However, this method of assessing wayfinding ability leads to the next section of this chapter that reviews a group of alternative, more objective methods of measuring wayfinding ability.

The simplest objective measure for assessing wayfinding performance is to record the amount of time taken to reach the destination[9]. This was used, for example, by Butler in (Butler, Acquino et al. 1993), who timed subjects' journeys to the nearest second, up until the moment they reached their wayfinding goal. This method was also used by O'Neill, in (O'Neill 1991), where he measured three variables, one of which was the amount of time spent reaching the wayfinding goal. Another factor he measured was the number of instances of 'backtracking' (turning round and re-tracing part of a journey). Finally, O'Neill also assesses the number of wrong turns made at choice points (which he defines as "a turn in a direction that is incongruent with the most efficient completion of the wayfinding task."). O'Neill's estimation of wrong turns is similar to Peponis' calculation of 'redundancy' as a measure of wayfinding difficulty in (Peponis, Zimring et al. 1990). This is a measure of the number of choice nodes passed through in excess of the minimum number to get from origin to wayfinding goal. These kinds of measures contrast significantly with the more subjective measures described earlier in this chapter.

Finally, through an examination of objective measures of wayfinding ability, Golledge develops and uses most tests in his paper (Golledge 1995). In this paper he conducts a wayfinding experiment using thirty-two subjects in a familiar, real environment. The subjects were instructed to make their way from an origin to a specified destination and then back again. Their journey times were recorded whilst making a note of whether the subjects took the same route from their destination back to the origin again. He then asked each subject to gauge (in a questionnaire) why they had chosen the route that they taken (they chose from a list of criteria); he then compared their routes to a number of pre-computed routes (based on a sample of the same criteria). Some of the criteria he uses are "Shortest Distance", "Fewest Turns/ Many Turns", "Longest Leg First/Shortest Leg First". Although the act of asking the students to select which criteria *they* thought most influential is a quite subjective test, Golledge also scores their routes objectively against the set of calculated routes.

He notices that for one route, 62.9% of subjects took the same route to and from the specified locations, whereas for an alternative pairing of locations only 15.6% of subjects took the same route each time. He suggests that this difference can be reasonably explained through the use of strategies that minimise either distance or the number of turns; however he suggests that this cannot completely account for the asymmetry of routes. He states, "This implies that, in addition to the previously discovered asymmetry of distance perception… perceptions of the configuration of the environment itself (particularly different perspectives as one changes direction) may influence route choice. Thus a route that seems shorter or quicker or straighter from one end may not be so perceived from the other end, thus inducing a change of route." The conclusion of his work is that the main criteria influencing route choice decisions are shortest path, fewer numbers of turns and selection of longest-leg first. However, what is particularly interesting about this result is the fact that the direction of the route appeared to influence the route chosen. ('Shortest path' and 'fewest number of turns' are variables that are direction-independent, only 'longest leg first' can be affected by direction.) Explanations of this phenomenon do not support the theory that a cognitive map is crucial to path selection. Clearly the path selected was affected by a set of visual variables *dependant* upon direction of the routes (whereas cognitive maps are independent of direction). Two of the questions Golledge then highlights for future research are "For non-route retrace what were the criteria that caused a different route choice for the return trip? What difference does it make to predicting travel when one uses different route selection criteria for outbound and inbound trips?" Both of these are particularly provocative questions, suggesting potential areas of future research.

In the next section of this chapter, instead of considering the wayfinding performance of people, ways to assess the impact of the world are considered. These methods too, fall into two categories: ways of appraising the environment in a subjective manner, and techniques for analysing an environment objectively. Gärling suggests three methods of assessing environments. Firstly, he uses the subjects' judgements of the degree of differentiation of the environment (do different parts of an environment look different?). Then he gauges the visual access (visibility of key locations) of the space, and finally considers the complexity of its spatial layout (he has difficulty defining this measure). He proposes a system for classifying environments based upon these three criteria; however, the assessment of each criterion is relatively subjective and he fails to fully define how each of these criteria should be judged.

In a manner similar to Gärling, Weisman has a concept of how environments could be analysed, yet is unable to execute it in a truly objective manner. In (Weisman 1981), he particularly considers plan configuration. However, he ultimately assesses this by asking a group of subjects to rate a series of stylised plans against four criteria - preference, complexity/simplicity, 'describability' and judged ease of wayfinding. In other words, in failing to create a set of methods for assessing plan configuration, he seeks the subjective opinion of his subjects.

Butler et al. also use subjects to assess their routes subjectively in (Butler, Acquino et al. 1993). After conducting the experiments, their subjects were asked to rate a variety of routes, on a scale of zero to ten, ten being an "optimal" route and zero being a "terrible" route. In their results, shorter routes seemed preferable (when tested against 'number of turns' or 'number of decision points'). One factor that may have some bearing on his results is the fact that, particularly in his second experiment, the subjects were being asked to make judgements purely on the basis of examining plans of single routes through theoretical buildings. If subjects had actually to navigate through the routes, rather than simply examining at them on plan, would their assessment have been the same? Nevertheless, irrespective of whether it is a real building or a plan of a building being judged, it is still a highly subjective way of assessing the complexity of an environment. The fact that Gärling, Weisman and Butler all fall into the same trap indicates that an objective analysis of the complexity of environments is not an easy task.

Some people have begun to approach this problem in a more objective manner, for example O'Neill. In (O'Neill 1991), he devises a topological method for assessing the complexity of a building plan, using a graph-based analysis where each decision point is represented as a node in a graph. The degree of complexity is a relatively simple comparison of links to nodes (in effect calculating an overall link-to-node ratio). Although a good starting point, it does have one obvious flaw, which is that by simply looking at the average distribution of choices available, he fails to take into account any methods of analysing the visual properties of the environment.

Consider the following example. Figure 2.2 represents the corridor system within two buildings. The first building consists of two, long, parallel corridors linked at regular intervals by five cross-routes, typically forming a series of T-junctions. The second building is similar to the first, except that the rightmost corridor has been replaced by a series of shorter corridors, at varying angles to each other. Try conducting the following thought-experiment. Imagine a person standing in the bottom-rightmost corner of each building. In the first building this person would have an unobstructed view down the longest length of the building and would be able to note the locations of the adjoining side-corridors. A person standing in a similar location in the second building would have no such view, as the visual information available would be far more limited. It could be argued that the second building is more complex, and it is feasible that it could cause more wayfinding problems than the first building.

Using O'Neill's measure of complexity (based upon choice-point nodes and links), both buildings would be judged as equally complex. This is because they contain the same number of nodes and links. However, when axial analysis is used to represent the buildings, this type of analysis is based upon the fewest and longest lines of sight that pass through every space in a building. Because it uses sight lines, this representation picks up the fact that visibility in the second building is reduced. This factor is indicated through the overall mean depth of the building. In the first building, this measure is 1.53 and in the second building, it is 1.80. The increased depth (higher number) of the second building implies greater complexity. The comparison of O'Neill's measure and axial analysis is shown in Figure 2.2.

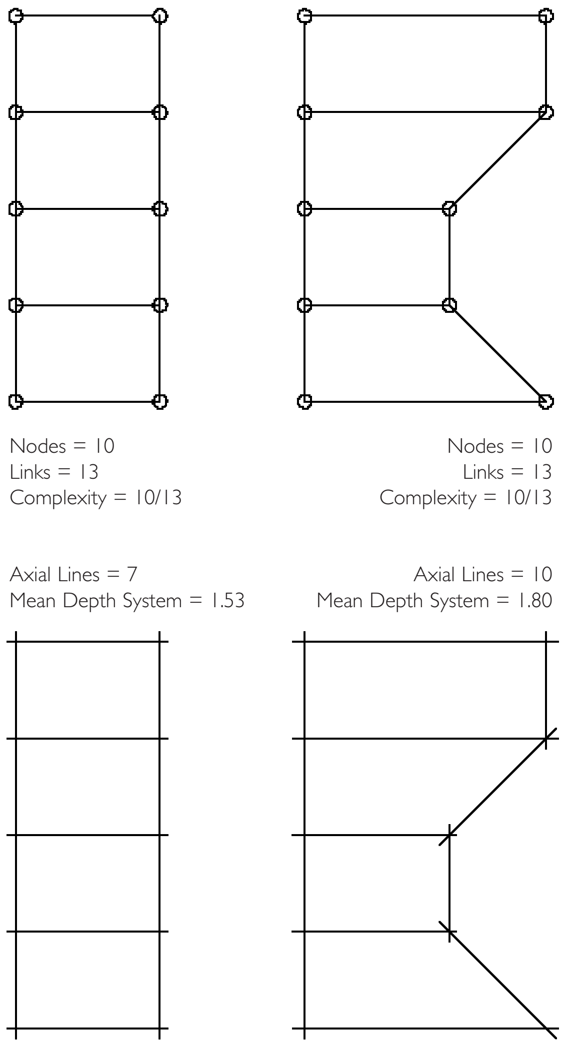

Figure 2.2 'Complexity' and Mean Depth of a System

O'Neill does, however, also make the point that, "other than these [papers by Braaksma and Peponis that he discusses] and a handful of other studies that attempt to objectively measure the environment, little research exists to relate the psychological and physical aspects of wayfinding to a larger conceptual model of legibility." Incidentally, although his method for assessing complexity was very simple, the results of his experiments were interesting. He reports that as the topological complexity of a plan increases, then the ease with which wayfinding tasks are performed decreases. It should be noted that O'Neill is one of the few authors who has attempted to use *both* an objective measure of an environment and an objective measure of subjects' wayfinding ability.

In the quote above, O'Neill mentions two other authors whose approach he admires, the first of which is Braaksma. Braaksma invented an ingenious way of analysing the complexity of an environment based upon the construction of a graph-based representation of a series of mutual visibility conditions. Braaksma's method starts by identifying a number of essential locations in an environment. In the example given in his paper, the environment being analysed is an airport, meaning that the types of location that Braaksma selects are, for example, the building entrance, check-in desks, passport control etc. He then constructs a large matrix, and tests to identify the locations that are visible from each of the other locations. He performs this test for each possible pairing of locations. Braaksma then goes on to develop a visibility index, which is a single measure of the overall, average visibility of key locations in the whole building. For example if everywhere were visible from everywhere else, this number would be 100%. The method he uses is very elegant, simple and effective. He goes on to report a relationship between this measure and perceived wayfinding ease. However, since the ease of wayfinding was judged extremely subjectively through informal interviews, this finding is not as significant as it might have been. Unfortunately, Braaksma has never continued this work.

Finally, the other person whom O'Neill cites positively in his paper is Peponis. In (Peponis, Zimring et al. 1990), they use a configurational method of analysing the building, based upon the longest and fewest lines of sight passing through every space in the building. They construct a graph-based representation of the relationships between these lines of sight (the crossing of any two lines is represented as a link in the graph; the lines themselves are the nodes). They then calculate how accessible any line is, on average, from all other lines in the system. The importance of the configurational properties of an environment to wayfinding is stressed by Passini in (Passini 1992), where he says that "Although the architecture and the spatial configuration of a building generate the wayfinding problems people have to solve, they are also a wayfinding support system in that they contain the information necessary to solve the problem." And "The comprehension of the principle by which spaces are organised appears as the single most important factor in facilitating image formation of a building [rendering it intelligible]."

After analysing the environment in this objective manner, Peponis et al. go on to observe and analyse the performance of the subjects as objectively as possible. The subjects are given wayfinding tasks with set pairs of origins and destinations. As the subjects find their way from an origin to a destination, they are followed by an observer who marks their route on a map. One of the primary measures of route analysis that Peponis then applies is 'redundancy' (mentioned earlier in this chapter). All in all, Peponis is one of the few academics who has conducted research into wayfinding and who has been able to consider *both* the environment and the performance of their subjects in an objective a manner as possible. It should therefore be of no surprise that this paper has been particularly influential in the field of wayfinding research.

If there is a criticism to be levelled at Peponis et al., it is due to the methods he uses to make the observations of the subjects' routes. In previous work done by Peponis and other Space Syntax researchers (see Chapter 1), the pedestrians being observed were completely unaware of their observation. This was easy to achieve because the kind of pedestrian movement being observed was what is termed 'natural movement', that is to say movement from everywhere to everywhere. In Peponis' experiments, due to the need for the experiments to be controlled wayfinding experiments with pre-determined origins and destinations, the observations could only be conducted using willing participants. These subjects would have been aware that they were being followed and that their movements were being recorded on plan. Although this was as objective as possible a method Peponis could have devised (at the time), there still exists a risk that the *presence* of the observers could have effected the subject's wayfinding performance. This risk reflects a limitation of the technology available at the time.

Wayfinding in Virtual Worlds

Today, if Peponis' wayfinding experiment were to be performed it would be possible to automatically track the paths of the subjects, without them being aware of the presence of an observer (akin to having an ever-present invisible observer). This could be achieved by using virtual environments to conduct a series of controlled experiments. Indeed, in recent years a number of researchers have attempted to work in this way, investigating how the use of virtual simulations may aid our understanding of wayfinding and navigation in unfamiliar environments. Such work dates back to the early 1990's whereas Peponis' paper, (Peponis, Zimring et al. 1990) was published in 1990, just as research using virtual environments began. At this point computers became powerful enough to simulate large-scale worlds and to allow people to walk through these worlds, or *wayfind* through them, in real-time. At the same time, hardware became affordable enough to permit such experiments. Since the early 1990's, the growth of papers on this topic has been exponential. During this decade, computing power has followed Moore's Law[10] (processing power doubling every eighteen months for the same cost), Therefore, the growth in papers on virtual wayfinding could be regarded as being fed by the accessibility of computing power. The growth in research on this topic is such that in April 1998, there was a special edition of the journal Presence devoted to Navigation and Wayfinding in Virtual Environments, followed by a second Special Edition on the topic in December 1999. It would be fair to say that it is a small but nonetheless flourishing subject area.

A number of academics entering this new research area originate from psychology or geography and many were previously involved in wayfinding research in the real world. However, this group of people have been joined by a substantial group of researchers who come from a computer science background. The origins of the range of researchers in this area mirror their concerns. Whilst many focus upon how wayfinding in virtual worlds may aid our understanding of wayfinding in the real world, others are beginning to see a necessity to understand the process of wayfinding and navigation in the virtual world, in order to design virtual worlds that are intelligible and easily traversed. This need becomes more crucial as the virtual worlds become more abstracted from the real world (for example in the case of three-dimensional visualisations of large datasets). As Darken says in his introduction to the 1999 special edition of the journal Presence, (Darken, Allard et al. 1999), "Few things are as fundamental to the human experience as the interaction between humans and their environment - be it physical or virtual." Later, in this same editorial, Darken and Allard go on to stress the dual nature of virtual wayfinding research by stating that, "In our attempt to make better interfaces for virtual environments, we must understand what carries over from the real world to the virtual world. On the other hand, in some cases, we want to go in the other direction: we want to carry skills or knowledge acquired in a virtual world to the real world." This duality of research aims and potential applications is also mirrored by Ruddle et al., in (Ruddle, Payne et al. 1998), when they say, "research should address the navigation of VEs per se as well as the transfer of spatial knowledge learned in VEs to the real world." The different goals prompting research in this area ensure a diverse focus of interest.

In order to begin to use virtual environments for wayfinding and navigational research it is first necessary to understand the technological limitations of the methodology. One issue that must be addressed early on in this body of research is to determine to what extent the technology *itself* may effect the outcome of the experiment. Only then can an answer be given to the question of whether we can learn from the study of virtual environments how people will behave in real environments. Some of the technological issues to be investigated include fundamental issues of interface design, such as those addressed by Ruddle et al. in (Ruddle, Randall et al. 1996). In this paper, they seek to determine whether there is a difference in patterns of movement between subjects navigating immersively and navigating non-immersively.

Ruddle compared the performance of subjects navigating through an immersive world, either immersively (using a head mounted display) or using a desktop (monitor) display. The virtual environments they navigated through consisted of two buildings containing an almost equal number of rooms. Whilst in the virtual worlds, the subjects performed direction and distance tests to gauge their spatial knowledge. One of the first (and in the context of this thesis potentially interesting) findings was that the subjects navigating using a headset could be seen to be taking advantage of this interface by looking around more than the 'desk-top subjects'. The ease with which this can be done immersively, in a way that feels quite 'natural', may explain this phenomenon. It was also found that the 'desk-top group' spent more of their time stationary than the immersive group. Nevertheless Ruddle also notes that contrary to his expectations, the proprioceptive feedback for the group using the headset did not drastically improve their orientation-judging abilities. On the whole, the conclusions of his paper were that similar patterns of movement were found between the immersive and non-immersive group.

Another important factor that plays a part in designing wayfinding experiments is the method by which motion is controlled through the environment. It is this issue that is addressed by (Peterson, Wells et al. 1998). In this paper they are primarily interested in determining whether navigational control devices that use whole body movements (and orientation) are superior to alternative control devices such as joysticks, which are independent of body position. In their experiment, the subjects were required to repeatedly follow a series of path markers; then the accuracy of their route learning was tested in the same environment without the markers. In addition to route-learning tasks, they were also required to perform some direction orientation tasks. It was initially found that navigational accuracy was higher using the joystick, and that route learning was equivalent for both devices. When it came to the spatial orientation tasks, it was found that there was a negligible difference in simple environments. However, the body-controlled devices were found superior as environments grew more complex. This is not to say that body-independent control devices inhibited their spatial cognition of the world, as found in (Bowman 1999), rather it appears that the performance of body-controlled devices is merely superior. Other work undertaken into interface issues which look into the effect of input devices include (MacDonald and Vince 1994; Wann and Rushton 1994; Slater, Usoh et al. 1995; Chance, Gaunet et al. 1998; Bakker 1999).

Another area for investigation is the effect of the restricted field of view (FOV) of most headsets. This was investigated by (Péruch, May et al. 1997) who performed a number of homing tests using headsets of differing fields of view (40°, 60° and 80°), finding no marked difference in performance between groups. They conclude that the amount of simultaneous environmental information is not a determiner of our comprehension of the environments. This is in contrast to an earlier paper by (Alfano and Michel 1990), in which Alfano and Michel found that restrictions in FOV affect the performance of certain tasks including the ability to follow a winding path, to perform a hand-eye co-ordination task and to form a mental representation of an environment. Note though, that since Ruddle found that subjects wearing headsets were more likely to move their head to a greater degree, it could be that we simply compensate for a restricted FOV by looking around more. Ruddle also cites earlier, unpublished work, in his paper, (Ruddle, Payne et al. 1998), where he notes that subjects occasionally travelled past locations just outside their FOV. However, he goes on to estimate that this behaviour accounted for no more than 5% of all navigational errors. In this paper, Ruddle also notes no difference between subjects' behaviour who are wearing headsets with different FOV (45° and 90°). Other work has been conducted into the importance of peripheral vision to motion perception, but clearly if virtual worlds are to be used for wayfinding research the precise effects of the various technical aspects of the experiment such as FOV need to be known and understood. This holds true for all other aspects of setting up experiments, whether they are desktop or immersive, and independent of the manner in which movement is controlled. In the context of this thesis, with regard to setting up any future experiments, the following guidelines can be applied. It is advisable, if using an immersive system, to have as wide a FOV as possible and to ensure that the physical orientation of the user contributes to their movement control (see Chapter 3).

Before continuing to look at the series of papers that describe various virtual wayfinding experiments, there is one other paper that is of relevance, which considers the methodological rather than the technical issues of setting up wayfinding experiments in virtual worlds. In the paper (Bowman 1999), Bowman considers how what he terms "travel techniques" can affect the spatial orientation of his experiment's subjects. The experiment was conducted in three travel 'modes'. In the first mode a subject follows a path chosen by the computer, with no control over their motion, yet they can look around whilst in motion (similar to being a passenger in a car). The second mode is similar to the first, the only difference being that the subject defines their own path before moving through the world; however once in motion they cannot deviate from their chosen path. In the final mode, the subject has complete control over their route choices during travel (analogous to being the driver of the car as opposed to the passenger in the first example). Bowman concludes, "For tasks in which spatial orientation is especially important, it appears that a travel technique giving users complete control over their position… can produce high performance levels [of spatial orientation]."

The result of Bowman's paper, suggesting that the best results occur when subjects are given greatest freedom of choice, leads to a series of papers which attempt to consider a number of different wayfinding issues in virtual environments. For example, one aspect of wayfinding that is far easier to investigate in virtual worlds (compared to the real world) is the combined effect of plan complexity and familiarity. Since virtual worlds can be fictitious (as opposed simulating real worlds), it is possible to generate a series of test environments of varying complexity. (The issue of how to begin to measure complexity was addressed earlier in this chapter; see the sections on O'Neill, Peponis and Braaksma). Since the worlds have no basis in reality, it is possible to *guarantee* that subjects cannot be familiar with them, and therefore *degrees* of familiarity can be strictly controlled. Two such papers that investigate this aspect of wayfinding are (Ruddle, Payne et al. 1998) and (O'Neill 1992). In O'Neill's paper, he examines the effect of plan complexity on wayfinding in simulated buildings. He finds that wayfinding performance decreases in proportion to an increase in environment complexity. However, as familiarity with an environment increases, the effect of plan-complexity is reduced. One explanation for this phenomenon is that perhaps we simply do not 'wayfind' in environments with which we are quite familiar. Therefore, a pre-requisite for the act of wayfinding should be that there is some degree of doubt regarding the correct route through an environment, and that once a route becomes well-known, wayfinding no longer takes place.

The most interesting research finding of this paper is evidence of an effect of plan complexity of spatial orientation. Environment complexity was measured using a definition developed in an earlier paper by O'Neill in (O'Neill 1991), using a measurement of the mean number of paths leading from a choice point in the world. O'Neill terms this measurement ICD or Interconnection Density (correct graph terminology would be 'mean node degree'). It was found that plan complexity (using O'Neill's definition) affected wayfinding ability; the more complex the building the longer it took to find the goal, and the more wrong turns were made by the subjects. By monitoring and comparing wayfinding performance against increased familiarity with the building, it was noted that the effects of plan complexity on wayfinding performance decreased with increased familiarity of the building. All of these experiments were undertaken in a simulated environment, based on the assumption that the results would be representative of performance in a real environment. In Ruddle's paper on the effects of familiarity, (Ruddle, Payne et al. 1998), his findings concurred with O'Neill's, namely that with increased familiarity of the virtual environment, subjects' wayfinding performance also increased. The means Ruddle and Payne use to assess wayfinding ability were: distance time to goal, and distance/orientation measures.

In Magliano's paper, (Magliano, Cohen et al. 1995), they consider the impact of the wayfinding goal on wayfinding performance. This experiment was conducted within a virtual simulation (in this case a simulation comprising a series of still images forming a walkthrough a small town). The extent of subjects' spatial-knowledge acquisition was measured after observing the pictorial walkthrough. Before the experiment, different instructions were given to subgroups of the participants. They were instructed to: either attend to possible landmarks, to the route itself, or to the spatial configuration of the small town. A control group were given no such instructions. Subjects' spatial knowledge was assessed by varying tests, which were applied to each group. Magliano et al. concluded that the subjects who had been given the instructions to learn only the landmarks did no better in the memory tasks than the route or configuration learning groups. The main distinction between the groups occurred during a direction-giving task. The subjects who had attended to either route or configurational information gave better (more accurate information) than either the control group or the landmark group.

This result suggests the relative unimportance of landmarks, leading to a set of papers, all of which use virtual environments to assess the impact of landmarks on wayfinding, with widely differing results. Lynch originally suggested that landmarks played a significant role in our cognition of the environment, in (Lynch 1960). Magliano's paper above seemed to suggest that for environment learning, landmarks had no effect upon a subject's spatial knowledge. This lack of evidence for any landmark effect is supported in the next two papers. The first of these is (Tlauka and Wilson 1994).

Tlauka reported upon an experiment in which subjects were required to learn their way through a sequential series of rooms (linked by two doors, one always 'locked' the other 'unlocked'), in a virtual simulated environment. For one group, memorable landmarks were placed in the environment to aid navigation by providing visual cues. For a second group, no such landmarks were present. No noticeable difference in task performance was found between the two groups. This unexpected result was explained by Tlauka and Wilson, in terms of strategy. They hypothesised that in the landmark case, room-landmark pairings were learnt, while in the non-landmark situation, sequences of right/left choices were memorised. They maintained that landmarks do contribute to navigation, but as only one of many navigational strategies used. They judged that the true effect of landmarks upon navigation can only be accounted for by sufficiently suppressing all other strategies or techniques. They then attempt to do this in an additional experiment (in which the R/L sequence-learning strategy is suppressed through the imposition of a simultaneous counting-backwards task). Again, there was no measurable effect of landmarks upon performance. Despite a result at odds with their expectations, they do, however, praise the utility of computer-simulated environments in the accurate testing of such navigational aids.

In Ruddle's paper (Ruddle, Payne et al. 1997), they investigate the effects of landmarks on route-learning ability and other spatial cognition tasks. They find a slight improvement in the time taken to complete the task in the environment containing landmarks in contrast to the environment without landmarks. However, when performing distance and orientation estimates, the effect of the inclusion of landmarks in the environment appears to be negligible. In contrast to this outcome that only weakly suggests that landmarks play any role in wayfinding, the subjects report in questionnaires that they actively used the landmarks, particularly in forming associations with specific locations in the world. The weak (as opposed to significant) effect of the landmarks appears to support the findings in Tlauka and Wilson and Magliano's papers.

In (Darken and Sibert 1993) and (Darken and Sibert 1996) the effects of landmarks on time taken to reach a wayfinding goal and on distance/orientation estimates were small. However, of all the differing navigational aids used in their experiments (landmarks being one such aid), some were of more use than others. Although it was not found that local landmarks significantly aided navigation, the inclusion of a virtual 'sun' as a global landmark did appear to improve performance. This distinction between local and global landmarks is subsequently addressed in a paper by Steck and Mallot (Steck and Mallot 2000), which provides the strongest evidence for the importance of landmarks to navigation.

Steck and Mallot's paper puts forward compelling evidence for the use of landmarks and describes a pair of particularly well-constructed experiments designed to investigate the dual effect of local and global landmarks. One of the techniques employed in this paper was to alter the relationship between the local and global landmarks, after the subjects had already navigated once through the environment, hence creating a conflict of cues. They concluded that some subjects used local landmarks, others global landmarks and others a combination of the two, whilst some people alternated between using local and global cues. The overall conclusion of this paper is that there appeared to be evidence that landmarks were being used by people when finding their way. Since this contrasts with earlier work (which shows only a weak effect of landmarks), it may well be that landmarks *do* play a role. It is likely that this is only one type of environmental cue used when navigating. However, Steck and Mallot also conclude their paper by endorsing the usefulness of virtual environments for this kind of research, particularly for the ability to create "conflicting environments", i.e. non-realistic environments - in this case, worlds in which environmental cues shift between subsequent journeys. They say "virtual environments are a valuable tool for navigation experiments, both for consistent and inconsistent environments."

Another component of environments hypothesised by Lynch to be crucial to wayfinding is the path (the full set being paths, edges, landmarks, nodes and districts). The use of paths in virtual worlds is investigated by Darken and Sibert in (Darken and Sibert 1996). In this study, subjects were required to perform a naïve search task in a large-scale environment representing open sea and islands. The subjects were searching for targets (ships) and the experiments took place in five different virtual (sea/island) worlds. Each task in each world was attempted under a number of conditions (with the aid of a grid, a map and both a grid and map). Darken and Sibert conclude that disorientation arises from a lack of directional, visual cues. Another unexpected observation they make is that they surmise that path-following is a natural human spatial behavioural characteristic, to such an extent that even when an *explicit* path is not evident, other environmental features such as coastlines or gridlines were used as implicit paths. This same experiment is described in more detail in (Darken and Sibert 1996), where the effects of grids (pseudo-paths) on wayfinding is that they appear to significantly improve performance. They surmise that grids are useful for providing useful orientation/directional cues.

These papers' focus upon attempts to find an empirical justification for the environmental components identified by Lynch, have led a number of researchers to attempt to use these principles in the design of virtual worlds. As stated in the beginning of this section, the dual nature of virtual wayfinding research has resulted in it concentrating upon two different applications: to inform our knowledge of wayfinding in the real world, and to design easily navigable virtual worlds. This next set of papers briefly examines research that has addressed issues of wayfinding in virtual worlds, but with the aim of improving *virtual* environment design. A large proportion of these papers have directly implemented Lynch's concepts of paths, edges, landmarks, nodes and districts as design principles for virtual world design.

For example, one such an application of Lynch-inspired design principles is the approach taken in a series of papers by Ingram and Benford (Ingram and Benford 1995; Ingram and Benford 1995; Ingram, Bowers et al. 1996; Ingram and Benford 1996; Ingram 1997). In these papers, the authors are particularly concerned with the design of abstract data-spaces and how to make them easily navigable. They examine methods to insert paths, edges, landmarks, nodes and districts into their world designs. Using their "LEADS" system, districts, landmarks and edges are computed from the spatial distribution of the data, they claim to aspire to being able to evolve paths from the movement of the users over time. Nodes are then formed by the intersection of paths. In the absence of being able to achieve path evolution at the time of their writing, paths are instead inserted into the world using computed methods.

The main problem that they encountered with this approach concerned a conflict between traditional 'paths' and the six degrees of potential movement available in the type of environments they are using. However, on the whole, they found that subjects performed wayfinding tasks in less time with repeated exposure to the "LEADS"-enhanced environments compared to plain environments. The subjects also claimed to feel less disorientated in the "LEADS"-enhanced environments. In a later paper, they also discuss Space Syntax research (Ingram and

Benford 1995), which leads them to conclude that the "subtle inter-relations between access, lines of sight, navigability and probabilities of social encounter can be exploited in the implementation of suitably designed or evolved virtual villages, towns and cities… in this way city (etc) metaphors for virtual environments may produce gradients of accessibility for information" In other words, they speculate that the use of city-like environments may serve to be useful metaphors when designing navigable abstract worlds. Ingram and Benford finally express a desire to combine both Space Syntax and Lynch-inspired approaches to aid navigation in future work on virtual worlds design.

Another important paper is that of Charitos (Charitos 1997) in which he proposes using Lynch-like components to aid navigation in the virtual world. His paper begins by describing possible examples of types of virtual landmarks, signs, boundaries (edges), thresholds (edges), places, paths, intersections (nodes) and domains (districts). Unfortunately, unlike Ingram and Benford, he makes no attempt in this paper to test the effectiveness of these objects on wayfinding task performance. The approach described in this paper remains conceptual only.

In (Darken and Sibert 1996), Darken also attempts to use some of Lynch's principles in order to investigate their effect upon wayfinding (this is in addition to his earlier work on landmarks and grids). In particular he suggests dividing the environment into smaller parts (districts) and ordering these parts using an organisational principle, such as a road network or an underlying grid. The conclusion reached by Darken was that "the presence of the wayfinding augmentations did significantly improve searching performance." although as a device, this was less useful than the provision of a virtual map. He ends this paper by noting that "Although not all wayfinding augmentations are appropriate for every problem, this research begins to show what types of information are most important, how they can be provided, and how they might be used. This is the first major step toward a methodology for designing navigable virtual worlds."

Finally, since it is felt appropriate by some researchers to apply real-world design techniques to virtual worlds to aid wayfinding, and that researchers attempting this approach appear to have met with some initial success, it is worth asking what work has been conducted into assessing the similarities between wayfinding behaviours in the real and virtual worlds. The last section of this literature review will examine what is known about the degree to which our behaviour in these two types of environment is analogous.

If we consider this final selection of papers in chronological order, it can be noted that the first paper dates from 1982. Although the experiments are not strictly conducted within a virtual environment (namely a *computer-generated* environment through which it is possible to walk in real time), this paper does compare the similarities between knowledge gained in a simulation compared to real-world knowledge. The fact that these kinds of comparisons were already being made prior to virtual environment wayfinding research is indicative of the interest in this research question. If the question underpinning this thesis is whether we can learn from the study of virtual environments how people will behave in real environments, then it is vital to understand what is already known about the similarities and differences between our behaviour in the two realms.

In (Goldin and Thorndyke 1982) Goldin and Thorndyke compare the transferability of spatial knowledge acquired in a simulated environment (in this case a film, rather than a true virtual world) to the real world. The simulated environment was a film of a route through West Los Angeles, shot from the inside of an automobile. The real tour was experienced as a bus-ride along the same route. Since it is accepted that driving through an environment will usually result in a differing pattern of environment-knowledge than walking through that same environment, then the direct relevance of this paper to this thesis could be held to be questionable. However, if the knowledge acquired in a simulated environment (albeit a driven-passenger simulation) is directly applicable to its real-world equivalent, then it may be that this transference can also be applied to different types of navigation (walking). This paper also prompts the question of whether 'walking' through an immersive virtual world is more analogous to walking in the real world, or to some other real-world situation. For example, it could actually be that walking through an immersive world is more similar to driving in the real world. This may of course be due more to motion perception and proprioception than to the visual experience.

Goldin and Thorndyke base their paper upon the theory that knowledge derived from direct navigation is superior to knowledge acquired from maps. The experiment by Goldin and Thorndyke divided their subjects into two groups: one group navigating the real environment, the other the simulated environment. (Sub-groups of each main group had their experience of the route augmented through either verbal narration or through a ten-minute map inspection). Each group was then tested for three different types of spatial knowledge; landmark knowledge, procedural knowledge (knowledge of a route) and survey knowledge (map-like knowledge). Goldin and Thorndyke predicted that there would be no difference in landmark knowledge, that the film groups should have less procedural knowledge than the real group and that the real-world group would demonstrate survey knowledge (map-like, spatial configuration knowledge) that was equal to or superior to the simulation group.

In contrast to their predictions, the simulation-group performed better on the landmark knowledge tests than the real-world group. This result could be attributed to their methodology; at every landmark, the film paused and zoomed in upon the landmark for ten seconds, making them potentially more significant than they might have been otherwise. In terms of procedural knowledge, there was no significant difference in distance and sequence estimations although the real-world group performed better for the orientation tests (approximately ten degrees more accurate in their estimations). To the surprise of Goldin and Thorndyke, the supplementary information (narration and map) did not appear to contribute to procedural knowledge. Finally, for the survey knowledge tests, there was no difference between the real group and the simulation group, excluding those members of the simulation group who also had access to a map (this subgroup out performed every other group). Although surprised by many of their results, their overall conclusions supported their hypothesis, namely, that simulations of environments can act as adequate substitutes for real-world spatial knowledge learning. The only situation where this may not be valid is where directional information is critical. For orientation tasks the simulation proved to be a less than satisfactory learning environment. However, this might not be true in an immersive world where the act of looking around is performed by a physical turning of the head or often of the whole body. Such physical body movements may help to orient a person in a virtual environment in a manner similar to the real world.

After this early paper, the remaining authors conduct their research in truly virtual environments. The first of these is (Witmer, Bailey et al. 1996), who examine how route knowledge gained in a simulation of a complex office building can be seen to aid navigation in the real building. They compare the performance of three groups; one group learns the route in the real building, the second in a virtual simulation of the building and finally one group learns the route solely from colour photographs. The subjects' spatial knowledge acquisition was measured initially by performing distance and orientation estimates. Their wayfinding ability (and hence route knowledge) was measured by recording the total time and distance travelled (using a pedometer) by each person and calculating their number of wrong turns (incorrect choices at an intersection, entering of wrong rooms and backtracking). The traversal time was almost equal for both the real and virtually trained groups, with the virtual group making slightly more wrong turns than the real group. Witmer et al. also estimate that configurational knowledge was unaffected by the method of training, and go on to conclude that "These results suggest that VEs that adequately represent real world complexity can be effective training media for learning complex routes in buildings."

Written in the same year as Witmer's paper was a paper by Tlauka and Wilson, (Tlauka and Wilson 1996). In this paper they conduct an experiment to test the spatial knowledge gained in a virtual world, compared to the knowledge gained through examining a map of the same environment. In particular they were primarily interested to see whether the knowledge gained in the virtual world was orientation-free - namely that the knowledge was flexible and independent of the orientation of the observer. Whilst instructing subjects to conduct orientation-pointing tasks, their time taken to perform the task was also measured. (The assumption being that if a subject has an orientation-specific knowledge of the world, extra time is required to mentally 'rotate' the map, before they can indicate a direction). The subjects' knowledge was tested through orientation estimates and a map-drawing task. The conclusion of this experiment was that the group that had studied the map had an orientation-specific knowledge of the world, whereas the group that had navigated through the environment had an orientation-free knowledge of the test-environment. Although this was not a direct real/virtual comparison, the important fact was that these results were similar to results found when comparing map-learnt environments to real environments. This caused Tlauka and Wilson to conclude that "the present study suggests that real-world and simulated navigation both result in similar (i.e., orientation-free) cognitive maps." They go on to say that previous work has shown that "there is a great deal of equivalence of learning in simulated and real space."

In the following year, there was a paper written by Ruddle (Ruddle, Payne et al. 1997), in which he reproduced an earlier study conducted in the real world by Thorndyke and Hayes-Roth (Thorndyke and Hayes-Roth 1982). The experiment, which had been originally constructed in a real environment, was reproduced in a desktop (non-immersive) VR. It was found that the users effectively learnt the spatial layout of the world, in a manner that was analogous to Thorndyke and Hayes-Roth's original experiment. Ruddle set three types of wayfinding test, route-finding ability (distance and time), relative distance and direction estimates. In both the real and virtual experiments, the participants were required to make direction and distance judgements. The results of these direction and distance judgements were found to be comparable with Thorndyke and Hayes-Roth's original experiment, thus enabling Ruddle to conclude that navigation in real and virtual worlds was comparable.

Finally, a series of papers written by Darken between 1997 and 1999 are all concerned with this issue of virtual-to-real knowledge transfer. (Darken and Banker 1998) (Goerger, Darken et al. 1998) and (Darken and Goerger 1999). In (Darken and Banker 1998), they concluded that exposure to a virtual simulation subsequently improved wayfinding performance in the real world. This particular experiment was conducted in a natural rather than man-made environment, and the task being performed was an orienteering[11] task. One group rehearsed the route using maps, the other using a virtual simulation. The performance of the subjects was monitored in a number of ways, and their actual paths through the environment were measured using a GPS[12] system. The result that Darken and Banker found by conducting this experiment was that a subject's level of prior experience of orienteering (beginner, intermediate or advanced) appeared to make the greatest difference to the experiment outcome. However, Darken ultimately concluded that a virtual environment was a useful training environment and that it had an added advantage which was that it was quicker (and less tiring) to train in the simulation compared to the real world.

In another study made by Goerger and Darken et al. in the same year, they performed a similar experiment in a complex building. This time they found, in contrast to all of the studies outlined above, the group that studied only a floor plan of the building performed *more* effectively than the group which trained in a virtual simulation. In the face of evidence from other papers (in particular the work by Witmer), they are unable to conclude that virtual environments are of no use for spatial knowledge transfer. However, they conjecture that their seemingly contrary results were due to the reduced time spent training in the virtual environment compared to the relative complexity of the environment. It may be that the more complex an environment becomes, the more time is needed familiarising oneself with the world.

Waller also stresses the time spent in a virtual simulation as being key to spatial knowledge-acquisition. In his paper (Waller, Hunt et al. 1998) he compares the wayfinding performance of subjects trained using a map, using the real world, a desk-top virtual simulation or an immersive virtual simulation (with both long and short exposure times). He concludes that, when only a short time period is spent in the virtual simulation, there is no advantage gained over using a map. However, with sufficient time spent in the virtual world, subjects can out-perform those trained in the real world. They conclude that, "With a few caveats, VEs can be an effective medium in which to train spatial knowledge."

Other researchers in the field broadly support the conclusion of Waller. Below are gathered a selection of quotations that effectively summarise what researchers in the field of virtual wayfinding are saying about the relationship between navigation in the real world:

- In (Darken and Sibert 1993) they conclude that "principles extracted from real world navigation… can be seen to apply in virtual environments."
- In (Witmer, Bailey et al. 1996), they state that, "These results suggest that VEs that adequately represent real world complexity can be effective training media for learning complex routes in buildings."
- In (Tlauka and Wilson 1996) they conclude that "navigation in computer-simulated space and real space lead to similar kinds of spatial knowledge."

However, a couple of notes of caution are voiced;

- "With a few caveats, VEs can be an effective medium in which to train spatial knowledge." (Waller, Hunt et al. 1998)
- "We need to better understand how spatial knowledge is acquired." (Goerger, Darken et al. 1998)

All of the above authors appear to be suggesting that we use real space and virtual space analogously, on the basis that knowledge gained in either one may be applied to the other. This result clearly has many applications, such as training people to navigate through environments that they are unable to use in the real world for training purposes. This can be because the real environment is hazardous (fire-fighting simulations or other emergency scenarios), or inaccessible (space simulations, hostage rescue situations) or does not (yet) exist in the real world.

However, the assumption that we navigate through real space in a manner that is analogous to virtual space based on evidence of knowledge transfer (between realms), is a fundamentally flawed assumption. Although the similarity of spatial knowledge gained in these two realms certainly supports the notion that real and virtual behaviour is analogous, by itself it is not sufficient evidence. In order to evaluate whether this technology is a viable technology to research wayfinding and navigational behaviour in the real world, it is necessary to determine whether we actually *use* space in the same way in both domains. To answer the question of whether we can learn how people will behave in real environments from the study of virtual environments it is vital to understand the similarities and differences between our behaviour in the two realms. It may be that we do indeed use space in subtly different ways, without this affecting our overall spatial comprehension of the environment. However, the assumption currently being made is that any cumulative effect (i.e. spatial knowledge acquisition) must always be a product of the same constituent acts (micro-scale behaviours). In fact, it could be possible that the cumulative effect (in this case, spatial knowledge) may arise from differing combinations of actions in either realm. The answer to the question 'Do we use space in a manner that is analogous and are our micro-scale behaviours and actions similar in either environment?' is vital to both wayfinding research in general and to this thesis in particular. It is also a question which has yet to be adequately addressed in existing research.

Chapter Summary

At the beginning of this chapter, wayfinding research in the real world was examined. The distinction was made between researchers focussing upon knowledge in the head (the wayfinding performance of subjects) and knowledge in the world (the design and layout of the environment). It was demonstrated that with few exceptions both the wayfinding performance of subjects and the analysis of the effect of the environments had been assessed using predominantly subjective methods. A case was put forward for the development of more objective ways of analysing both wayfinding performance and the role played by the environment. In particular it was felt that it was important to consider environment and behaviour together rather than in isolation. Methods developed by different researchers for analysing wayfinding performance and the layout of the building were presented.

Once research began into wayfinding in the virtual realm, researchers took a number of different approaches. Some felt it important to study the effect of the technology and interface, to determine how to best set up experiments in order to test wayfinding performance, and reduce any effects caused by experimental methods. Other researchers performed wayfinding experiments by simply substituting the virtual world for the real world. In particular a group of researchers used the ease of computer-generated *theoretical* environments to facilitate investigations into the effects of the environment, and in particular to test out Lynch's hypothesis that landmarks are necessary visual cues to aid wayfinding. Little conclusive evidence was found to support Lynch's ideas. However, a number of researchers went on to suggest how the inclusion of not only landmarks, but also paths, nodes, districts and edges could render virtual worlds more intelligible and hence prevent disorientation in large-scale (and especially abstract) virtual worlds. Finally, a number of researchers investigated whether spatial knowledge gained in a virtual simulation of an environment could be usefully applied to the real world. The assumption made, was that if the knowledge gained in both types of environments was comparable, then wayfinding behaviour must be analogous across the realms. An argument is put forward for why this is an inadequate method for comparing real and virtual navigational behaviour. The importance of seeking answers to this, in the context of both wayfinding research (in general) and this thesis (in particular), is stressed. In order for the work of this thesis to be relevant (to either virtual navigation or real world pedestrian movement research), it is necessary to determine if the two are comparable. In the next chapter, one method for assessing this relationship will be presented, as it is felt vital to establish this fact before any further experiments can be conducted.

Key Points

- This thesis' definition of wayfinding is that it is the act of travelling to a destination by a continuous, recursive process of making route-choices whilst evaluating previous spatial decisions against constant cognition of the environment.
- Methods of measuring wayfinding ability or behaviour are inadequately developed and suffer from being subjective.
- Methods of analysing aspects of the environment and gauging the environment's effect on wayfinding performance are inadequate and also suffer from being predominately subjective measures.
- There has been little research undertaken into directly relating wayfinding performance back to the design of the environment.
- In comparing wayfinding in the real and virtual world, it has been assumed that since spatial knowledge gained in a virtual simulation may be successfully applied to the real world, therefore wayfinding behaviour in both situations must be comparable. No work has been done to determine whether we navigate in a similar manner.

- Route asymmetry (taking one route from A to B and a different route from B to A) is an interesting phenomenon that may serve as a clue to investigating wayfinding research issues.

Notes

1 Earlier papers tend to use the word in its hyphenated form, whereas later papers tend to use the concatenated form.

2 From The Oxford Dictionary of English Etymology and An Etymological Dictionary of the English Language, 1924.

3 Source from The Oxford English Dictionary, Second Edition.

4 Source A Dictionary of Americanisms, Matthews.

5 Tetris is a simple computer puzzle game. As small shapes fall down the screen, they must be rotated to fit together to complete lines. When an entire line is filled with blocks, it is removed from the screen. If the player cannot complete lines, the blocks will eventually rise past the top of the screen and the game ends. It was invented by Alexey Pajitov in 1985 whilst at the Computer Centre of the Academy of Sciences in Moscow.

6 Source, Raubal, M. and M. J. Egenhofer (1998). "Comparing the Complexity of Wayfinding Tasks in Built Environments." Environment and Planning B 25(6): 895-914.

7 Source, Raubal, M. and M. Worboys (1999). "A Formal Model of the Process of Wayfinding in Built Environments." Lecture Notes in Computer Science(1661): 381-400.

8 Source, Raubal, M., M. J. Egenhofer, et al. (1997). "Structuring Space with Image Schemata: Wayfinding in Airports as a Case Study - Spatial Information Theory - A Theoretical Basis for GIS, International Conference COSIT '97." Lecture Notes in Computer Science 1329: 85-102.

9 Although there remains a slight problem caused by people walking at different speeds, one solution being Witmer's use of a pedometer in (Witmer, Bailey et al. 1996).

10 "The observation that the logic density of silicon integrated circuits has closely followed the curve (bits per square inch) = 2^(t - 1962) where t is time in years; that is, the amount of information storable on a given amount of silicon has roughly doubled every year since the technology was invented. This relation, first uttered in 1964 by semiconductor engineer Gordon Moore (who co-founded Intel four years later) held until the late 1970s, at which point the doubling period slowed to 18 months. The doubling period remained at that value through time of writing (late 1999)." Source, The Jargon Dictionary, http://www.netmeg.net/jargon/terms/m/Moore_s_Law.html.

11 A "Competitive sport in which runners cross open country with a map, compass, etc." Source Pocket Oxford Dictionary (electronic version), Oxford University Press, 1994.

12 Global Positioning Satellite. A system whereby a user's location is determined by triangulating their distances from a number of satellites in stationary orbit. The accuracy of different systems varies.