UNIVERSITY OF LONDON

SPATIAL NAVIGATION IN IMMERSIVE VIRTUAL ENVIRONMENTS

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE BUILT ENVIRONMENT

UNIVERSITY COLLEGE LONDON

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

DEPARTMENT OF ARCHITECTURE.

BY

RUTH CONROY

LONDON

JANUARY 2001

To Sheep

Abstract

Pedestrian movement studies in real environments have shown consistent statistical relationships between 'configurational' properties of spatial layouts and movement flows, facilitating prediction of movement from designs. However, these studies are at an aggregate level and say nothing about how individuals make the micro-scale decisions producing these emergent regularities. They do not therefore 'explain' movement. Progress on this is difficult since decision-making mechanisms are hard to observe in the real world and the 'experimenter effect' is ever-present.

Could the study of movement in immersive virtual environments help? If it could be shown that movement in virtual environments was analogous to movement in real environments then micro-behaviour data (head movement, direction of gaze, visual search behaviours) could be obtained through virtual experiments. The aim of this thesis is to explore this possibility by constructing experimental worlds with spatial properties varied to reflect those known to relate to movement in the real world, and asking individuals navigate through them immersively.

Powerful analogies are initially demonstrated between virtual and real behaviour. Two types of micro-scale analysis are then performed: linear analysis, examining how routes are formed and how far linearity is conserved, using measurements of cumulative angular deviation along a path, string-matching algorithms to determine average routes, and analysis of isovist attributes along routes; and positional analysis, focussing upon pausing behaviour, including examining where subjects pause along routes, what choices are made at junctions, how isovist properties of pause-locations compare with an environment's overall isovist attribute distribution, and correlating pause-point and isovist data. In each analysis, 'subjective' movement behaviour is related to 'objective' properties of environments.

The experiments show results strongly suggesting how noted aggregate regularities are produced: linearity is strongly conserved, usually following long sight-lines, with pauses in configurationally 'integrated' locations offering strategic visual properties, long lines of sight, and large isovist areas.

Contents

Acknowledgments

I should like to thank Professor Bill Hillier for his advice and guidance throughout the course of this thesis. Thanks are also due to Professor Michael Batty for his comments on the final draft and to the many colleagues who have contributed to this research.

To the EPSRC for their generous funding.

To thank the Tate Gallery for permission to use their building plans as a basis from which to generate a virtual environment for these experiments and to the Space Syntax Laboratory for the use of their observation data for this building.

To the Institute of Contemporary Art (ICA) in London for inviting me to participate in the "Parallel Space: The Geography of Virtual Worlds" one day conference and permitting the conduction of an experiment upon their premises. To the delegates of the conference who participated in this experiment.

To all the subjects of all other experiments, mainly graduate students of the Architectural Association School of Architecture and the Bartlett School of Architecture, whose enthusiastic participation was invaluable.

To Stefan Ponur, who constructed the 3d model of Barnsbury, when I unable to finish the construction of the virtual worlds in time to conduct the experiments.

To Nick 'Sheep' Dalton and the MSc Course in Virtual Environments at the Bartlett for the extensive and unlimited use of their Virtual Reality headset and accompanying equipment.

To Nick 'Sheep' Dalton for the use of his analytic software, Axman, SpaceBox, Hyper Hyper Pesh and Infinity Within and for his assistance with the writing of OmniVista, the isovist analysis software written specifically for this thesis.

To Jenni Scott for proof-reading and commenting on the final draft.

To Christopher Woodward for his typographical advice.

I would particularly like to thank Joe Burlington for his support and guidance on managing my time effectively.

Finally, once again to Sheep (this time in his capacity as my husband) who originally encouraged me to undertake this work and who has supported me tirelessly over the years. If a spouse could be awarded an honourary doctorate for putting up with the trials and tribulations of doctoral research, he is fully deserving of one.

Illustrations

Illustrations Continued

Illustrations Continued

Tables

Chapter One:
Introduction – Goals of the Study

Architects, planners, geographers and other social scientists are currently unable to accurately predict pedestrian movement in the built environment. The need to make such predictions becomes particularly important when changes or interventions are being made to an existing urban environment. Equally, this knowledge gap can also be a problem when designing large, complex buildings (the larger or more complex a design problem is, the more probable it is that a designer's intuition will fail). In order to be able to make these kinds of predictions we need to strive to understand better how the built environment is *currently* being used by people. This is the approach that has been adopted by a body of work known as Space Syntax, an area of research that forms the foundation to this thesis.

Space Syntax Analytic Methods

Space Syntax is a family of theories and methodologies concerning the social use of space, which grew from the collaborative research of Hillier and Hanson at University College London in the mid-1970s. From the very beginning, Space Syntax research focussed on the relationship between space and social life, be this the social life of a simple building, a complex building (or set of buildings), a settlement or an urban district. At the time, the approach pioneered by Hillier and Hanson was unique, as it involved analysing solely the spaces between buildings as opposed to a building's (or buildings') geometric form (Hillier and Hanson 1984).

Through their research, it emerged that one vital component for the generation and maintenance of an environment's social life is pedestrian movement, since this type of movement (in contrast to, for example, vehicular movement) creates a potential for social interaction. This social interaction may be overt and direct (informal meetings around the photocopier in an office) or it may be indirect and discrete (the social policing of an urban area; personal safety afforded through the continual co-presence of strangers) (Hillier at al. 1992, Hillier 1996). Over the years, Space Syntax researchers have gradually come to the conclusion that pedestrian movement is one of the *best* indicators of the social success, or conversely the social failure, of an environment. This relationship between the environment and its inhabitants' through-movement lies at the heart of all Space Syntax research.

Gradually, researchers at UCL developed a series of new techniques for spatial analysis that can be used to *predict* pedestrian movement. Such prediction is possible due to the discovery of a statistically significant correlation between attributes of spatial configuration (of an environment) and pedestrian movement. One of the reasons why they were able to uncover this relationship stemmed from their focus upon space rather than geometric form. Space can be considered as not merely the 'gaps' left between buildings (or walls) but more importantly as *potential* for movement. If space is regarded in this way, then a strong and direct relationship between space and movement seems less surprising. However, over the years, one form of spatial analysis has proved to be consistently better at predicting pedestrian movement (and even vehicular movement) than all the others developed at UCL. This method of combined representation and analysis of space is termed 'axial analysis'.

Axial analysis consists of an evaluation of the spatial structure of a system by examining the primary lines of sight (or 'axial lines') that pass through it. The first stage of the analysis is one of representation. It is necessary to fully represent the entire spatial structure of an environment by means of axial lines; this process is termed creating an 'axial map'. The identification and recording of the longest and fewest lines of sight present (such that every space in the system has at least one line of sight passing through it) is a prerequisite of axial map production. Every line of sight will cross at least one another line (since only continuous systems can be analysed using Space Syntax), creating a network of axial lines that mutually intersect at intervals. It is this complex relationship of axial line intersections that forms the basis of the second stage of axial analysis.

In the second stage, each axial line is further represented as a node in a graph. When any two axial lines cross, this is indicated by a link between the two nodes representing those lines. Gradually a graph-based representation of the entire system can be constructed based upon nothing more than the intersection-relationship between axial lines. This graph is topological, containing no information about scale or distance. Once the graph representation has been built, the relative position of each axial line within the whole system can be assessed. This is achieved by calculating the number of steps in the graph (equivalent to real-world changes of direction) that it is necessary to follow, on average, to move from any line to all other lines in the system. Those axial lines that are, on average, a greater distance (in graph terms) from all other lines are said to be 'segregated' within the environment. Conversely, those lines that are only a short distance (in graph terms) from the rest of the system are termed 'integrated'. The reciprocal of this value (which is known as mean step depth[1]) is known as the 'integration value' of an axial line and it is this value that has consistently correlated well with mean pedestrian movement flow data. It should be noted that currently the best indicator of pedestrian movement is the measure termed 'radius 3 integration'. This measure is computed in an identical manner to the calculation for integration (as described above) except that it only considers lines that are three steps away. For this reason, it is also termed 'local integration'. Therefore, at the core of Space Syntax research is the notion of relationship, and in particular the relationship between each space (whether is it represented by an axial line or another means of representation) and every other space in the environment under analysis.

As Space Syntax's analytic methods have developed over the years, there has been an ongoing process of consistently checking all theoretical techniques against observed pedestrian movement patterns in the real world. This ever-increasing database of empirical data constitutes one of the largest collections of pedestrian movement observations and this empirical observation data has also served as an invaluable resource for this thesis. However, the majority of the observations in the database were taken at the level of people 'en masse' - cumulative, aggregate movement counts. Much of the data collected covers, for instance, the average flow of pedestrians along a street or road-segment and mean occupancy rates of rooms or spaces. Far fewer observations have taken place at the level of the individual. As a consequence of this, there currently exists a gulf of both data and knowledge between movement at the level of the individual and people's cumulative behaviour. It is not clearly understood which small-scale actions or decisions made by a person in an environment can result in the large-scale observations that form the Space Syntax database. However, does there exist a *need* to bridge this gap between large and small-scale movement? As Space Syntax analysis is *already* able to predict overall patterns of pedestrian movement should this not be sufficient for research purposes and the requirements of design professionals? Partly, this thesis is responding to the challenge that there exists a gap in current knowledge; the fact that the relationship between the individual and the population is so poorly comprehended is intrinsically an interesting research problem.

Pedestrian Movement and Emergent Phenomena

The fact that an enticing research area *can* be identified may be insufficient reason to undertake research; potential applications of the research should also be established. One reason to investigate how a person's actions may result in large-scale movement patterns is to understand better such cumulative behaviour. This is particularly necessary if overall pedestrian movement is considered to be an example of emergent phenomena. Can we even begin to categorise pedestrian movement as an emergent phenomenon if we have no knowledge of the micro-scale forces causing the large-scale movement?

For what reasons could pedestrian movement be regarded as an example of emergent phenomena? One, extremely clear definition of an emergent phenomenon is that it is an overall effect that stems from a small number of simple actions (or rules), *but* that the outcome *cannot* be easily predicted from the initial set of basic steps[2]. For example, using the above definition, patterns of axial line integration generated by Space Syntax analysis can *certainly* be regarded as an example of an emergent phenomenon. In this case, the 'set of simple rules' involves little more than determining whether any two axial lines are intersecting. The rule for generating the topological graph and then calculating the integration value of each line is also remarkably straightforward (refer to the earlier description in this chapter). However, the resultant pattern of integrated lines and segregated areas can be both surprising and revealing. It has even been known for quite experienced Space Syntax researchers to be unable to confidently predict the resultant integration pattern before performing the analysis. Since overall pedestrian movement mirrors the distribution of integration in an environment, it is extremely tempting to also describe this as an emergent phenomenon.

However, for movement patterns to be classed an emergent phenomenon, then it *should* be possible to determine the set of simple rules that give rise to these patterns. One example of where pedestrian movement can be seen to be an emergent phenomenon concerns the behaviour of tightly packed crowds. This type of behaviour is clearly demonstrated by the work of Benford et al. in (Benford, Greenhalgh et al. 1997). However, these results cannot be successfully scaled up to the urban level. It would appear that we behave quite differently when in a crowd situation, for at these times our actions are mostly governed by the movement of our neighbours. When moving through a city our behaviour is far less affected by the relative location of other people.

It is suggested that there may be a number of simple rules that give rise to natural movement in an urban environment. This premise formed the foundation of the paper (Mottram, Penn et al. 1999) in which programmable, software 'beings' (with vision) were given a set of simple instructions. Their instructions were to follow a long line of sight (with some minor random deviations) until a new, 'longer' line of sight is noticed and subsequently this new direction is followed. It was found that these rules gave rise to cumulative patterns of movement *suggestive* of real-world movement patterns. Given the tentative success of this earlier work, then it could be reasonably hoped that if a set of rules could be shown to give rise to overall movement, then this would provide the 'missing link' of Space Syntax research. It is not disputed by critics of Space Syntax that statistically significant correlations are found between observed pedestrian movement flows and patterns of integration; an explanation of the correlation is often sought and questioned. Why should we, as individuals, move in a manner consistent with the integration values of axial lines? This lack of a causal explanation has long been an 'Achilles' Heel' of Space Syntax research. So, not only would the establishment of a set of rules governing natural movement be able to confirm that urban-scale movement *is* an emergent phenomenon, but it could also begin to clarify why Space Syntax research is able to correlate pedestrian movement so strongly with space.

Observational Methods

If an identifiable need exists for such research, why has this need not been addressed before? Why, for example, have observations been predominantly made at the level of the population and not at the scale of the individual? If more observations had been made at individual levels, then our understanding of pedestrian movement as an emergent phenomenon might be more sophisticated than it currently is. At first, the focus upon cumulative movement patterns seems to be a curious imbalance, but in fact, there are a number of historical reasons why this has arisen.

Firstly, it is far easier to measure large-scale movement accurately. Minor, individual idiosyncrasies are 'smoothed' over a larger sample size, ensuring a consistently high statistical correlation (with spatial configuration). The act of making observations at the aggregate level is also far easier, since measures such as 'mean pedestrian flow per hour' rely purely on the accuracy of a quantitative record (how *many* people passed this point?)[3]. For movement observations where location data is necessary (such as directional movement flows), the relative inaccuracy of location estimates is often in keeping with the coarser resolution of spatial analyses of large-scale urban environments and therefore has never been perceived as a problem. In contrast to the relative ease of large-scale observation techniques, there are both technological and methodological problems associated with attempting to gather data at the level of the individual.

The various methods for observing individual movement can be categorised as manual (by direct observation) or automated (utilising technology to aid observation making). The family of methods that are used to manually observe individual pedestrian movement consists predominantly of the subject being discretely followed by an observer who marks the subject's route onto a plan. The subject need not be actually moving, as a variant of this method is used for making static observations (although in the case of static observations the location of multiple people will usually be noted on the plan). The first problem associated with this method of making observations is the inaccuracy inherent in the process of researchers transcribing subjects' locations onto a plan, since this relies on the researcher's *judgement* to estimate the subject's location relative to the environment. As this judgement is subjective then the result is clearly open to error. On the whole, up until now, this has not been an issue, since such fine-scale observations are rarely used, and as long as the *general* path taken by the subject though an environment is noted, this is usually sufficient for most research purposes.

Another problem associated with this method of data collection is a phenomenon termed the 'observer effect'. This occurs when the presence of the observer within the same environment as the subject affects the subject's behaviour. It is possible to attempt to avoid this through the discretion of the observer. However, even allowing for this, it is still extremely difficult to gauge the true effect of an observer's presence.

Finally, the last drawback associated with observing pedestrian movement manually, is the difficulty of recording sequences of different actions (taken by a subject). For example, it is certainly possible to follow, and to mark on a plan, a route taken by a subject. It could even be possible to approximately mark onto the plan any locations where that subject stopped or paused en route. It would be difficult, however, to try to accurately time the duration of their pauses, or note in which direction they are looking. It is not that it is *impossible* to perform this task manually, rather, it is the inherent difficulty in making multiple, simultaneous observations by hand.

Since there are a number of problems caused by making manual observations, one would think that it should be possible to make observations using mechanical or electronic means (i.e. to remove the observer from the process). Unfortunately, there are also problems associated with the automation of observing subjects in an environment. For example, although it is possible to use a video camera to initially record the movement (as in the Millennium Crowds Project, undertaken by Space Syntax Laboratory), computers are as yet unable to efficiently interpret this kind of data. The evaluation of this data still relies on an act of laborious (and potentially erroneous) manual interpretation. There are computer applications available that are capable of analysing video sequences, such as 'EthoVision' by Noldus Information Technology. EthoVision is a video tracking system that is capable of performing video tracking, track analysis and visualisation. However, this programme and similar applications usually rely on the location being observed being a 'closed system' (nobody leaving or entering the scene) coupled with prior identification of the objects (in this case persons) being tracked. This scenario rarely happens in buildings or urban environments. This software is better suited, therefore, to the tracking of rats in a maze.

It is possible to use a GPS[4] receiver to track people, in the way that Darken and Banker did in their paper (Darken and Banker 1998), in which they describe an orienteering experiment performed in a natural environment. They used GPS to record the subjects' paths as they traversed through the environment. However, even with current technology, the accuracy of satellite positioning can have an average margin of error of 13.2m in the United Kingdom[5]. In other words, this would probably be an even less accurate method of recording individual pedestrian movement than through making manual observations. Darken was only able to use this successfully as a method since he was investigating orienteering ability, and his subjects were traversing relatively large distances, in which an error of a few metres was negligible.

A popular and economical method of recording pedestrian flow is to use infrared beams coupled to electronic counters that are triggered whenever their beams have been interrupted. Examples of these are commonly used in many public buildings (such as museums and art galleries) to count visitors. Although this is accurate and (usually) unobtrusive, it is rare that enough of the devices are used to get an overall picture of movement. However, this method of recording movement is actually measuring the aggregate level movements, since it would be almost impossible to track an individual through such a system (unless they were alone in the building). In fact, the above method is actually measuring continuous pedestrian flow.

It is possible to accurately locate individuals within a building using an 'Active Badge' system. The Active Badge was conceived, designed and prototyped between 1989 and 1992 at AT&T Laboratories. An Active Badge is a small device worn by each inhabitant of the building, which transmits a unique infrared signal every ten seconds. The building is fitted with a series of sensors, which are designed to detect these transmissions and relay them back to a central database. The location of the badge (and hence its wearer) can be calculated using the information provided by these sensors. The largest single system that is currently in use is at the Cambridge University Computer Laboratory, where over two hundred badges and three hundred sensors are employed. Although it is possible to visualise this data graphically on a plan of the building, there has been no attempt to use this data to evaluate the building's performance or to relate it back to the environment in any way. However, this is a system that could be ideally used to track individuals through an environment. Perhaps in the future, when such a system is widely available and affordable, it will be used for such purposes.

One method that *can* be used to investigate individual movement utilises the nascent technology of virtual environments. Why might virtual environments provide the very key to understanding how the small-scale actions of people result in pedestrian movement? There are two primary advantages that such technologies bring to bear upon the problem of individual movement. Firstly, when using a virtual environment to 'observe' pedestrian actions, it *is* possible to make extremely high-resolution observations. Not only can the location of a subject in an environment be accurately measured (as well as the direction in which they are facing if the set-up is an immersive system), but the rate of measurement can be as frequent as every tenth of a second. In this way, every small step, hesitation, change of direction or brief glance can be recorded in the fullest of detail. The kinds of observations that it would be extremely valuable to make in the real world suddenly become possible to make in virtual worlds.

If observations can be made in this way, it is hypothesied that it *could* be possible to identify certain actions or behaviours, which, when repeated over the duration of an entire journey could be seen to produce a pattern consistent with either patterns of observed real-world movement or with patterns of integration (or both). It is suggested by this thesis that the key to gaining insight into this problem is empirical data and that virtual worlds can be used as a way to amass a large quantity of high quality empirical data.

Another advantage of using virtual environments can be seen by examining the concept of the 'scientific method'. One aspect of the scientific method concerns the procedure of conducting experiments. It is advised to attempt to reduce the number of variables present in an experiment, in order to properly observe the effect of the one variable being tested. Ideally, a researcher should strive to eliminate all extraneous variables (from the experiment) leaving only the *single* variable under scrutiny. In the case of real-world pedestrian movement, the number of variables that may be affecting or influencing route choice are too numerous to list. Not only are there too many real world stimuli to take into account but there may even be effects of the environment that we are, as yet, completely unable to identify.

Not only are real environments intrinsically complex, but there is an additional factor that serves to heighten the complexity of the problem, namely the effect of other people. Although, as stated previously, the presence of other people has a lesser effect on 'natural movement' than in a crowd scenario, this is not to say that the presence of others has *no* effect upon movement. One hypothesis[6] is that we are more inclined to select a well-populated route (which would result in a 'multiplier effect'). On the other hand, under certain situations (an ill-lit street at night, for example) certain people might be disinclined to walk along a street devoid of others.

In a virtual environment, there exists the advantage that the creator of the world (or the designer of the experiment) knows precisely what components went into the creation that world. A researcher using a virtual test-environment can be confident that the world will consist of only those factors that the researcher has chosen to include in the world. Therefore, the factors that might effect pedestrian movement in a virtual world are both finite and *potentially* knowable[7]. Variables being investigated (for their effect) can be examined by including or altering *only* those variables in the world. For example, it could be possible to create worlds that contain as little as space and form. It would be possible to test the effect of building appearance on movement by varying the textures incorporated into the world. The effects of street lighting or signage could be tested and even the presence of others. The presence of others could be generated by populating the world with avatars (people-like, programmed automata) to gauge any resultant change in movement. In essence, virtual environments constitute the ideal test-bed since the experimenter has complete control over the environment in a manner that is entirely impossible in the real world.

The advantages of using virtual worlds to gather empirical data can only exist if these results bear some relation to the real world. The fundamental question of this thesis is, therefore, whether we *can* learn from virtual environments about behaviour in the real world. Are real and virtual movement analogous in any meaningful or useful way? How far do human behaviours in simulated, virtual worlds reflect equivalent behaviours in the real world? If it could be established that it is possible to learn about real world behaviour from observing the actions of subjects navigating through virtual simulations, then this could serve to break open the observational-bottleneck currently impeding research in this area. With the ability to make a large number of observations over a wide variety of simulated worlds (each designed to specifically test one spatial variable) it could ultimately be possible to provide an answer to the question of how large-scale movement arises. From this answer, it should also be possible to authoritatively state why Space Syntax can be used as a predictor of movement (i.e. why there exists a positive correlation between pedestrian movement and patterns of axial integration).

Outline of Proposed Research

This thesis will begin by taking an environment for which the Space Syntax Laboratory's observational database contains high quality, meticulously observed movement data. The first stage of the experiment programme will be to reproduce these observations in a virtual world to determine whether the two (real and virtual) are, in any way, analogous. This process can begin to establish whether we can learn about real-world behaviour from the virtual.

This initial experiment may serve to highlight some of the technical and methodological challenges of this paper. One question that must initially be broached is to determine how easily virtual worlds can be used for this purpose and whether appropriate and relevant data *can* be gathered. An answer to this question must be found in the course of this thesis. For example, will a wide range of people *be able* to navigate through these worlds with ease? If such methods of generating data can only be used with subjects who are already sufficiently familiar with the technology, then the results will not represent an adequate population sample. The usability of the technology is, therefore, an issue that must be addressed early in this thesis.

The very design of the virtual worlds might also serve to be an interesting research challenge. For example, one question that may arise in the course of this thesis concerns the kinds of worlds that may be created and the limitations of the technology. The primary limitations of artificial environment construction are computer memory and processing power. In an immersive system, an over-large model (in memory terms) may produce problems for the subjects navigating through it[8]. However, it is a requirement of this thesis to discover and work within the capabilities and limitations of this technology. Although this is not a primary aim of this thesis it is, nevertheless, a factor that will be of relevance throughout its course. Answers to this question should be held to be a useful research outcome.

The next stage of this thesis will be to design a series of worlds that can be used to further investigate small-scale behaviours and actions made by people whilst navigating through them. The tasks that will be performed in these environments will be a series of wayfinding tasks in which all subjects start from one location and attempt to navigate to a pre-defined destination. By giving all subjects the same task, it is possible to discern any variance between individuals in a sample (and more importantly, any similarities). Whilst people are navigating through the virtual worlds any actions taken or decisions made en route will automatically be recorded. This data can then be compared to the spatial and visual characteristics of the different environments. A full and detailed description of the experiments' set-up and procedure will be given and the data so collected will be presented.

Once the empirical data has been gathered, the routes taken by the subjects may be investigated. The first technique that will be brought to bear on this problem is string matching. The section of the thesis that comprises string matching is semi-methodological as a new application of this technique is being introduced. This section will start to investigate how small-scale actions produce overall patterns of pedestrian movement by assuming a 'top-down' approach, namely by examining entire routes. The string matching analysis presented in this section will then be used to seek patterns in a sample of different routes. It is a new method for comparing a number of different routes and determining what patterns lie therein. The search for patterns is central to understanding movement as an emergent phenomenon, as it is the overall patterns that characterise it as such. At the end of this section, all empirical data will be analysed using this technique and hence may reveal any patterns in the data.

The next section of the thesis will also be semi-methodological, as it will introduce a novel method for analysing the 'straightness' of a route. Once again, this is a 'top-down' approach since this technique considers the characteristics of the route or path as a whole before considering what local rules might have given rise to such patterns. After analysing all of the route-data from the experiments, it is hoped that the outcome may contribute to answering the question of what kind of small-scale actions are important. At the end of this section, a solely methodological section will be introduced which will examine the changing field of view along a route. Although no substantive conclusions are drawn from this technique, it serves to introduce the notion that what we look at, or what information is visually available to us as we move through an environment may be critical to determining our actions. After considering the visual field along the routes, the thesis will then shift its focus. Instead of considering movement from a 'top-down' perspective (considering whole routes) it will employ a 'bottom-up' technique (actions that take place along a route). It is hoped that this duel-approach will combine into one method to investigate how the aggregate patterns of observed movement arise from a set of small-scale decisions.

The final section of this thesis will begin by considering whether there are any locations in the environment that correspond to specific actions of the subjects. The first method used for examining small-scale actions will investigate where people pause for significant amounts of time. The basic movement data will be analysed and every location where a subject pauses will be recorded. In a manner similar to the route string matching approach, a technique for identifying trends in the data will be applied which will aim to seek out any patterns in the spatial distribution of the pause points. Equally, in a manner similar to the examination of the visual field along a route, the direction of a subject's gaze at the pause points will also be recorded and visualised. This section is both methodological and substantive.

After considering where people pause, an attempt will be made to investigate whether these locations relate to any attributes of the worlds' spatial configuration. In order to investigate such a relationship, techniques of isovist analysis will be expanded. A method will be used that generates isovists in a grid filling all navigable space throughout the world. Finally, the attributes of the visual field at the pause points are compared to the overall visual field to determine whether people are more likely to pause in some locations as opposed to others. The findings of this chapter are essential to the understanding of how small-scale behaviours produce pedestrian movement.

It is hoped that this sequence of examining the small-scale actions outlined above may begin to provide an insight into some of the questions concerning the nature of pedestrian movement raised earlier in this chapter. However, before any empirical work can be undertaken, it is first necessary to review research that has already been conducted in the area of navigation and wayfinding. Wayfinding, as an area of academic research, is of the utmost relevance to this thesis since this is an area that has *started from* what is known (or is knowable) about movement at the level of the individual. In this way it can been regarded as a 'bottom-up' process of investigation as opposed to the 'top-down' approach taken by researchers in areas such as Space Syntax. Therefore, the following chapter (Chapter 2) will examine past and current research into wayfinding in both the real world and in virtual worlds.

Conclusions

The methodology being proposed in this thesis begins by establishing to what degree real and virtual movement are analogous. It will then continue by setting up a series of experiments, which will take place in different virtual worlds designed specifically to elicit a broad range of subject responses. Methods must then be developed for comparing the routes taken to other routes (to determine individual similarities and differences) and to a variety of spatial and visual environment-attributes. It is hoped that it may be possible to begin to understand, particularly through the comparison of small-scale individual actions to space, the simple rules that result in the emergent phenomenon which is pedestrian movement. It should also be possible to formulate a theory about why overall patterns of pedestrian movement correlate so well with axial line integration.

Key Questions

- Can we learn from virtual environments how people will behave in the real world?
- Do patterns of small-scale, individual actions exist that can be seen to accumulate, producing observed, aggregate patterns of pedestrian movement?
- Is pedestrian movement an emergent phenomenon?
- Can the study of individual pedestrian movement aid our understanding of the relationship between movement and spatial configuration?
- Are pedestrian movement in the real world and patterns of virtual world navigation analogous?

Notes

[1] This is known as 'status' in graph theory.

[2] Turkle defines an emergent phenomenon as "one whose components parts interact with sufficient intricacy that they cannot be predicted by standard linear equations; so many variables are at work in the system that its overall behavior can only be understood as an emergent consequence of the myriad behaviors embedded within." Source (Turkle 1997).

3 For further details on observation methods see (Vaughan, Major et al. 1997).

4 GPS is an acronym standing for 'Global Positioning Satellite'. This is a system whereby a user's location is determined by triangulating their distances from a number of satellites in stationary orbit.

5 Source from the GNSS Flight Recorder Approval Committee.

6 This hypothesis arose from observations made in (Conroy 1996) in which eye tracking experiments were performed, during which subjects were required to look at urban scenes. It was noted that where such scenes contained people, that these people were heavily scrutinised by all subjects in preference to all other environmental factors.

7 Of course, there does exist the possibility that a factor, which is not expected to have an effect, may actually do so. In which case the creator of the virtual world may not even realise that is could be an effecting factor.

8 One problem that may be caused is 'time lag'. This is when there is a noticeable difference between head movement and image refresh rates causing a 'lag'. This effect is also thought to be a contributing factor to 'VR sickness'.