Chapter Nine:
Discussion and Conclusion

Abstract

This chapter takes the conclusions of Chapters 3 through to 8 and combines their individual conclusions whilst discussing their implications in a wider academic and application-based context. It concludes with a statement of future research directions that could arise from the work forming this thesis.

Without question, the primary finding of this thesis is that it is possible to use virtual environments to gather data about people's likely behaviour in the real world. In essence this data can be regarded as being 'free data' which is inherent to all virtual world simulations. All systems for navigating through virtual environments use this data to continually update the visualisation of the world, be it the world as represented through a sophisticated headset or as displayed upon a standard computer monitor. This data consists of: a subject's location within a scene, their orientation (which way they are facing), a description of which objects can be seen by the subject (which surfaces are occluding other surfaces and which are completely hidden) and finally the commands given by the user to move around and re-orient themselves within the environment. Without such data, it would be impossible to navigate a three-dimensional virtual world (although many users are unaware that this data exists), since it is kept 'behind the scenes' and hence can be regarded as transparent. It so happens, though, that this is *also* the very same data that can be used to research the behaviour of people within an environment. This is not a case of generating data since the data already exists. The methodological challenge is to find a way to tap into and to retrieve this data for further analysis.

Not only can a subject's behaviour be sampled quite rapidly (in the case of these experiments at a rate of ten times every second), this rate exceeds that at which we can easily observe similar behaviour in the real world. This shows one advantage of using virtual environments for this kind of research: the ease at which navigational data can be collected exceeds any real world equivalent. Not only can data be sampled more rapidly, but also more accurately.

The real-world observational data that was used in this thesis was gathered by hand; although we are beginning to see the use of automated techniques for investigating people's behaviour in the real word, such as the use of active badges, video capture/analysis, infrared beams and pressure pads, they are still techniques that are in their early stages of use. The majority of observations are still done by hand. The analogy that can be applied here is that the difference between real-world observations and virtual navigational data is akin to turning up the focus on a microscope; the resolution of the data is far finer for virtual world data, with less 'noise'. If it can be demonstrated that people's patterns of navigation in virtual worlds are analogous to real-world movement, then the use of virtual worlds to investigate navigation and other behaviour should grow in popularity.

Another advantage to the use of virtual environments to investigate patterns of navigation is the range of world types that can be used. The worlds that can be used in such experiments can range from simulations of real buildings or urban environments to entirely fictitious, theoretical worlds. For this reason it would be possible to design purely theoretical worlds which tested one particular spatial variable which a researcher was investigating, and hence to gauge the effect of that particular variable upon movement. If the essence of the scientific method is to reduce and isolate the possible number of variables, then this can be done using theoretic virtual environments in a manner that would be nearly impossible in the real world. The other advantage of using virtual worlds is reduction of noise in the data. Not only is this noise reduced through the automated method of the data collection but, the design of the worlds themselves can eliminate noise. Users can navigate through 'pure', 'exact' environments that are completely different to the visually complex environment that we inhabit daily.

It should be mentioned that an important fact concerning the potential future use of such techniques was that, in the experiments forming part of this thesis, subjects found it extremely easy to navigate through these environments immersively. The usually steep learning curve required to learn how to move through these worlds was in fact very shallow, making the technology ideal for such experiments. It meant that subjects could be used from all occupations, and did not need to be familiar with this type of technology. Ultimately this is an advantage since it allows for a more representative sample population. At the moment, with current technology it seems that it is far easier to navigate through an environment immersively than it is to navigate on a desktop.

The thesis begins by addressing what was probably its predominant question. The question that was first posed in Chapter 1, was whether the research findings obtained in virtual worlds could be held to be indicative of real world behaviours. The reason why this question is so essential to the thesis concerns future applications of the research. If these two patterns of movement can be regarded as being analogous then it may be inferred that all the findings from the rest of the thesis may be applied to real world pedestrian movement too. Without this correlation, the results may be intrinsically interesting, but the applications are extremely limited. The fact that this thesis held the answer to this question to be so important was also interesting since many researchers in this field assumed that this relationship could be taken for granted.

After comparing the movement patterns for the Tate Gallery it would appear that people move through virtual worlds in a manner that is analogous to the way in which we navigate in real environments. The correlation of the two sets of data for the Tate Gallery was found to be particularly good. Since the virtual simulation of the Tate Gallery was only a spatial model and contained no art or other people, such a high correlation was all the more surprising. This result implies that the building (and spaces) plays a far greater role in determining where people walk than any other single factor.

Other factors that effect the correlation of the data are the physical constraints of the worlds. It was assumed, as one of the basic requirements of this thesis, that the physical constraints of the virtual world should emulate real world constraints in order for this to be used as a viable research tool. Navigation in virtual environments can be quite different from the real world. In virtual environments, it may be possible to fly, to pass through walls or even to jump instantly from one location to another. For this technology to serve as a useful research environment these abilities need to be curtailed, such that the experience is as close to walking in the real world. The good correlation yielded by the Tate Gallery data supports this decision to create a virtual environment that also simulated solidity and gravity.

Having determined that we *can* learn from virtual behaviour, how people are likely to act in the real world, the thesis continued by conducting a series of wayfinding experiments in six, distinct virtual worlds. It became clear at the end of Chapter 4 (when all of the movement paths for all worlds were presented together), that even from this simplest of visualisations, i.e. prior to any additional analysis, that route path data can be extremely informative with regard to how people were using space. It was concluded at the end of Chapter 4, that people appeared to be moving through space in a manner consistent with the chapter's accompanying Space Syntax analyses. Observations that people appeared to be moving syntactically were made based on the relative locations of the most and least popular routes.

It became evident that a measure of route popularity could serve as a useful tool to analyse the subjects' route paths. It was found that it was possible to compare routes by means of a cumulative measure of the spaces through which a subject had passed. This was achieved by representing each space in an environment by an ASCII text character and then representing any route through an environment as an ASCII string. Once a route could be represented as a string, it was possible to perform string matching analysis on any two routes to determine now similar they were to one another.

String matching analysis is a common method that has many varied applications. The first part of using this technique involved nothing more than applying an existing technique to novel problem. The second stage that was employed could be seen as an innovative modification of string matching techniques. Instead of merely comparing pairs of routes to determine how similar they were, it was decided to compare each route to every other route in the sample. Eventually a single number, termed the MNLD (mean, normalised Levenshtein distance) value could be calculated for each route. This number would indicate how similar that route was to every other route in the sample. The routes that were most similar to all other routes could then be regarded as being the most representative routes of the sample. The routes that were most dissimilar could be held to be the most idiosyncratic of all the routes.

Once again, a similar method to the technique for visualising routes by their angularity value could also be applied to visualising routes by this new similarity measure. When visualising routes by their mean angularity, the values of the sample of routes were matched to a colour spectrum, the straightest routes being coloured red and the most undulating routes coloured blue. Since every route could also now be assigned a similarity measure too, the most representative routes of the sample could be coloured red, whilst the more idiosyncratic routes could be shaded blue. These two methods of route visualisation could be usefully applied in conjunction. In this manner it could be determined at a glance whether the most popular routes were the straightest (as is the case) and that the least popular routes were the most meandering.

It was noted that the most popular routes in a sample were frequently the 'straightest'. This observation led naturally to question what kinds of route choice decisions were being made, such that they produced this phenomenon. The most obvious type of location to use in this attempt to analyse subject's decisions were road junctions, as these are locations where a route choice decision *has* to be made. The paths made by all subjects through one world were broken down to the level of their constituent road segments and junctions. The decisions made by each subject at each junction were painstakingly recorded. It was determined that any single route through an environment could be assigned a score based upon the culmination of the decisions taken at each junction where it was necessary for the subject to make a route choice decision.

The method of assigning a score to a route is achieved by noting, at each junction along a route, the range of possible route choices available and comparing these to the choices actually made by the subject. By noting the choice made at each junction, the journey of a subject could be expressed solely in terms of the decisions made by the subject. Equally, all the decisions made by the subject could be summed to produce a cumulative score. If the average decision (in degrees) made by a subject (for a single journey) is plotted alongside the maximum, mean and average angles at each junction (also averaged over the route), then the choices made by subjects appear to lie closer to the maximum angles than to either the average or the minimum. It was also noted that the variance and standard deviation of the 'angular difference' between the chosen angle and the available angles for all 306 junction-decisions in this sample is less for the angular difference between the chosen angle and the maximum angles than for the mean and minimum angles. In other words, it appears that subjects are choosing the straightest possible routes as opposed to the more meandering or undulating routes. This is a particularly significant result that supports hypotheses made by Hillier stating that people tend to follow the longest line of sight that approximates their heading. This finding would certainly begin to suggest the type of micro-scale decisions necessary to produce the aggregate patterns of pedestrian movement observed by Space Syntax researchers, over the last twenty years at UCL, that correlate so well with axial analysis. By following the longest lines of sight, a subject is both deviating as little as possible and behaving axially.

One interpretation of why people should wish to move in as straight a line as possible concerns human memory and complexity. It could be suggested that the act of deviating as little as possible serves to reduce complexity in an otherwise extremely complex environment. There is a well-documented phenomenon termed route angularity, which is mentioned in the following papers (Tolman 1938; Sadalla and Montello 1989; Montello 1991). Route angularity is the effect of judging a route that contains many changes of direction to be longer than a straighter route of identical length. This might also be linked to the 'magic number' in psychology. This was a finding by Miller in (Miller 1956) that stated that people found it easy to remember (short term) up to seven items, give or take two. These two findings begin to suggest why people unconsciously attempt to steer a straight path. It would be an interesting area of future research to investigate the effects of route angularity using virtual worlds, and in particular to determine whether or not 'seven, plus or minus two' junctions or changes of direction were significant.

The overall conclusion of this section of the thesis was that the choices made by subjects at junctions appeared not to be randomly made. This finding can be linked to the previous conclusion formed by analysing the isovist properties of pause point locations. Since it has already been suggested that people are not pausing randomly (in terms of location), but strategically. It can then be further suggested that there is also a pattern to the kinds of decisions that people are making as well. Note, however, that both of these findings appear to be the result of unconscious behaviour. In the questionnaires given to the subjects to complete after participating in the experiments, there was a question asking whether subjects had used any specific strategy to aid their wayfinding task. None of the answers given could give rise to the results described above.

In an attempt to further analyse the results of the route angularity finding, the absolute angle selected at any node was plotted against the maximum angle of incidence for each node. Regression analysis was performed and the resultant value for r-squared was 0.342. This implies that factors other than angle of incidence contribute to route choice decisions. It is suggested that the other factor determining route choice decisions is approximate direction or heading. Namely, that subject will choose the greatest angle of incidence at a junction on condition that it is in the approximate direction that they are heading. Once again, this entirely supports the theories put forward by Hillier.

Finally, as a way of concluding this section of work, a method was proposed to visualise the results of the route angularity analysis. Since the cumulative value of all decisions made along a route could be calculated it was also possible to determine the average decision made at any junction. Routes could then be ranked in order of the average angle turned through during the whole journey and subsequently colour-coded accordingly. This technique appeared to provide a valuable method of visualisation, supporting an intuitive estimation of a route's 'straightness'. Whilst still considering micro-scale behaviours along routes, a representation of the changing visual field along a route is presented. This representation is termed the route vision profile.

The thesis then shifts its emphasis from entire routes (linear behaviours) to smaller-scale actions taking place at instance along a route (positional behaviours). One kind of micro-scale behaviour that it was possible to calculate from the raw data, was the location in the world where people were pausing. This could be represented as a 'point' on plan, superimposed upon the original route path. The first conclusion to be drawn from this method of visualisation was that people appeared to be pausing at or in close proximity to road junctions. Another way of describing this phenomenon was that people appeared to be pausing in locations where a route choice decision needed to be made. However, people were not pausing at every junction that they encountered, there appeared to be only a few locations where they paused. One interpretation of this was that subjects appeared to be pausing at locations where a route choice decision needed to be made, but where the correct or appropriate choice was not immediately obvious. Where the decision *was* immediately clear, the subjects appeared not to pause or hesitate at all, but simply continued with their journey. It was only where the subjects were unclear of the choice that they needed to make that they paused to look around and to scrutinise the environment to aid them in their decision making. Another conclusion arising from this section of the thesis was that the combination of direction of gaze and pause point analysis appeared to be the most useful method of combining the data available in a graphical form. Once this data had been combined, it was not only easy to determine where a subject had travelled, but where they had stopped for long periods of time and while they had been stationary, in which direction they had been looking in order to inform their navigational decisions. The most intriguing aspect of this extremely simple way of visualising the data, was the way in which it was hard not to attempt to interpret and recreate the 'story' of an individual journey.

After the calculation and visualisation of pause points the use of a clustering algorithm was investigated, to determine whether or not this proved to be a useful statistical method. The method that was used was the k-means clustering algorithm and it was determined that it could be used effectively to identify groups or clusters of pause points. It was felt that although the use of this method did appear to achieve its goal, that is to identify any clusters of points, it was not felt that it performed any better than a researchers intuition. It rarely revealed patterns that would not have been evident simply by examining a visualisation of the points. Moreover, it was felt that since the algorithm did not take into account the spatial layout of the environment it could sometimes attribute a number of points to a cluster, when clearly they bore no relation to one another because they were spatially separated (perhaps by a visual barrier such as a wall). For this reason, it was concluded that although it had been a useful tool to investigate it really contributed little to the thesis as a whole. However, for future applications, where much larger data sets may need to be analysed over a larger number of worlds, then such an automatic method of cluster analysis could prove useful. Ultimately were it possible to combine some sort of cluster analysis alongside a form of spatial analysis, then this could prove to be an even more useful method of analysis. For example, it could be proposed that the first test of whether or not a point should be assigned to a cluster is if it is visible to all other points in the cluster. This would not only combine a variant of cluster analysis but also include the types of intervisibility relationship discussed in the previous chapter, Chapter 8.

A quite significant result of the chapter on pause points was the fact that a comparison of total journey time or distance travelled to stopping patterns may reveal individual characteristics of journeys. This was never developed further since it was not felt to be central to this thesis. However, it appeared, that it could be possible to discern characteristics of a number of journeys, indicating possible strategies of different subgroups of people. In the experiment presented in the thesis it was suggested that the results could be showing both gender, age and left/right-hand biased differences in navigational patterns. Without question, this could certainly be an avenue for future research, since gender differences in navigation strategies and abilities are often discussed and do form the subject of certain psychological investigations.

Almost through an accident, another possible research finding from this thesis became evident. Based on data from two of the seven worlds in this thesis, it could be suggested that there is a possible correlation between the speed of a subject’s movement and the duration of their pause points. The faster a person moves, the less time they are likely to spend stationary. This finding came about, when towards the end of a series of experiments, subjects were shown how to move faster (the full explanation for why the experiment was conducted in this manner can be found on page 66). The result of this increase in pace appeared to produce a hitherto unanticipated decrease in the duration of pause points. It was not that people were pausing any less frequently, it seemed rather that people were merely pausing for shorter amounts of time. It was as if their judgement of time was modified through their speed of movement. This could without doubt be a factor worth future investigation, particularly when researching navigational patterns in virtual worlds as a stand-alone research area. This could be regarded as a distinct area of research in opposition to this thesis, whose focus is to investigate virtual world navigation in the context of the light that it may shed upon pedestrian movement in the real world. There may be future applications for the results of research such as this by designers of virtual worlds. Even today, there are designers who are being commissioned to design virtual worlds that need to function appropriately, virtual worlds that *are* the design product (rather than an interim stage to an ultimate realisation in the real world). For example the work being done on collaborative virtual environment at the University of Nottingham or at the BT Research Laboratory. These are worlds that are inhabited and support a range of activities and social interactions. If it is verified that a subject may pause less frequently or for reduced periods of time, depending upon their pace, then this could have a direct effect upon the design of these virtual worlds.

Another set of conclusions that arise from the detailed examination of pause point locations concern the differences between urban areas and buildings. It was surmised in Chapter 7, that there may be possible differences between stopping behaviours in a building and at the urban environment level. These differences may be investigated through analysis of the duration of pause points. It appears that people are pausing more frequently in a building than in an urban system. One possible conclusion that can be drawn from this is that a time interval that constitutes a significant pause at the building level may not constitute a significant pause at the urban level. It may either be the scale of the environment that is producing a difference in stopping patterns or it may be that there is a more rapid rate of optic flow (to use Gibson's terminology) in a building, the visual stimuli are changing more rapidly and therefore there is a need to pause more often to reassess these changes. Once again, this could provide a useful research question prompting new research directions. The use of virtual simulations of urban districts and buildings to determine differences at the micro-scale level between the different environments.

It is also possible that people pause in different types of locations in buildings and urban systems. In an urban environment, people appear to be pausing in proximity to road junctions. Any analogous behaviour is hard to discern at the level at the building. However, if the results of the Tate Gallery are compared to the locations of pause points in urban environments then it can be seen that people are pausing towards the centres of spaces (areas of below average isovist drift). These locations could be regarded as optimal locations for making route choice decisions in a building, since in the centre of a room, it is possible to look around as see through all doorways into adjacent rooms. It is less likely that a route choice decision is made at a door threshold in a building, whereas due to the difference in scale, it is entirely plausible that such a decision be made at a junction in an urban environment. For this reason, it may be that the finding that people pause in the centres of spaces in buildings is a manifestation of the same effect of subjects pausing at junctions in urban environments.

After determining that people appeared to be pausing in, or in close proximity to road junctions, it seemed that it could be a useful exercise to determine whether or not road junctions could be characterised through attributes of isovists generated at these locations. The conclusion of this analysis was that junctions appeared to be able to be characterised as areas of high mean isovist radial length. For certain environments, other measures may also be good indicators of junction locations; mean radial length, however, appeared to be a consistently good indicator over the sample of seven worlds. It is not that junctions are characterised by the areas of highest mean radial length, but rather there is a consistent identifiable increase in this measure at junctions, when compared to adjacent streets.

The observation that people appeared to be stopping at junctions and that these may be areas of high isovist mean radial length, led to an attempt to develop a typology of junction type based on its isovist's properties. This was proposed firstly as a useful method for examining people's behaviour; people's decisions at junctions can be examined in the context of the junction type (see Appendix B for a description of these proposed typologies). It can also be regarded as a fulfilment of Benedikt's prediction that isovists could be used in this way, namely to characterise different types of spatial arrangement.

Upon suggesting that it should be possible to develop an enriched typology of road junctions, it was then suggested that it should be possible to take a real world junction-isovist and fit it to one of the idealised junction-types (also see Appendix B). One possible application for this technique could be the automated mapping of road configurations. It should be possible to generate a series of isovists (either randomly or on a grid) such that all navigable space contains at least one isovist viewpoint (and ideally a number of points). The isovist resolution, as expressed by the number of isovists per square metre, or the grid divisions, can be varied. Once the areas of suspected junctions are identified (for example through the attribute high mean radial length), the individual isovists can then be matched to the idealised junction types. Once a match is confirmed, the junction location can be stored as well as the number of roads forming it and their angular distribution. The technique that proved most useful for achieving this goal was the use of the mathematical technique known as canonicalisation. Canonicalisation essentially describes the technique of matching any real-world junction to a canonical (or idealised) example. An example of how this technique might be applied to junction types was successfully illustrated by a real world example. By using such an automated environment mapping technique, an environment could be represented as a simple graph representation. Each junction could be represented as a node in a graph and every street (or uninterrupted line of sight) represented as a link in the graph. This suggested application of the technique of junction identification to environment mapping, is very similar to the method used by Kuipers, in his paper (Kuipers and Byun 1991) except that Kuipers used SONAR instead of visual fields, and his mapping robots were mapping real environments. The potential for an automated environment-mapping tool is high, and for this reason, this should be regarded as a potential area of research stemming from this thesis.

To complete the body of work conducted with junctions in urban environments, it was queried (in Appendix B) whether or not the shape of an isovist as generated at a road junction could be described by a mathematical equation. The conclusion was that junction shapes could be described by equations by using a polar graph and generating variations of the absolute value of the reciprocal of the cosine of theta. Ultimately it could be possible not only to automatically identify the locations of junctions in an environment, but also to record their exact shape as a mathematical equation. In this respect, an urban environment could effectively be described in an extremely succinct and economical manner. It could be possible to store descriptions of many urban environments extremely efficiently. However, to achieve this goal further research into fitting mathematical equations to junctions would need to be undertaken.

As a method for further linking single isovist properties to Space Syntax research, it was demonstrated statistically that certain attributes of the geometric shape of isovists correlated extremely well to syntactic measures of isovists. The implication of this is quite significant. It implies that it should be possible to infer properties of the configuration of a whole system from the visual field of the single isovist in which an observer is standing. This direct connection between the global and the local properties of isovists may confirm long held 'hunches' by Space Syntax researchers. However, a confirmation of the ability to make inferences about the interrelationship of spaces from nothing more than the single isovist that is immediately perceivable, is a finding that is quite remarkable. The very fact that there is also a correlation between pause points locations and specific isovist data only serves to reinforce that this is indeed a significant finding.

The connection between isovists and pause points was determined towards the end of Chapter 8 (the previous chapter). At an earlier stage in this thesis, it had been observed that people appeared to be pausing in or in close proximity to road junctions. However, this observation was purely subjective, prompted by the method of visualising pause points as dots on a plan. In this chapter, however, this observation was tested statistically. Both the central limit theorem and the z-test were applied to the task of determining how representative of the population of grid isovists (flooding all navigable space with isovist viewpoints) was the sample of isovists generated at pause point locations. The conclusion of this exercise was that people appear not to be pausing randomly in terms of isovist properties. It could be seen that on the whole, subjects were pausing in locations that offered strategic visual properties. These were locations that afforded unusually long lines of sight and large isovist areas. They were isovist locations that were highly integrated within the system, as well as highly connected and finally it appeared that people were pausing in locations that were far from any occluding surfaces such as building edges. The overall conclusion was that people were being exceedingly strategic in terms of where they stopped to scrutinise the environment. They paused only in locations that offered the maximum visual, local and global information about an environment, reducing the necessity to pause more frequently. The manner in which people navigated through an unfamiliar environment could therefore not only be seen as being strategic but also maximally efficient. It was demonstrated how such actions could cumulate into aggregate patterns of movement consistent with patterns of spatial integration.

This thesis has focussed on the question of whether it is possible to learn from behaviours in virtual environments how people are likely to behave in the real world. This chapter, in turn, will make some suggestions, not only for how this research might be applied, but also put forward some ideas for future research directions.

Future research directions and applications

Firstly, it is essential that the work of this thesis be continued as only seven worlds were used in the experiments, which is a relatively small number, and only two of these worlds were simulations of real environments. Although this thesis provided an excellent starting point for such research by suggesting that real and virtual patterns of navigation were analogous, this body of data represents only initial research steps. It is highly likely that the relationship between the two patterns is not a direct one. It may be that people's behaviour in a virtual world does differ somewhat from real world pedestrian movement. The precise nature of this relationship should continue to be explored. It would be unfortunate if the conclusions of this thesis were to be taken at face value and further work were to be undertaken simply assuming that one implied the other.

One way in which to undertake continued research into the relationship between real and pedestrian movement could be quite simply achieved. Every time a real world building or urban environment were analysed and pedestrian movement recorded (either as pure research or as consultancy) then an equivalent, virtual environment experiment of the same building or urban area could be conducted. If virtual world experiments shadowed real world observations, then over time a database of equivalent movement traces could be established. This would assist greatly the understanding of the relationship between pedestrian movement in the real world and patterns of virtual navigation.

There is another method that could be used to gather virtual movement data. Ruddle conducted experiments demonstrating no discernible difference between how people navigate in desktop virtual environments and immersive worlds (Ruddle, Randall et al. 1996). A future scenario can be imagined where a model of a proposed building is put online, for example in VRML format. The intended, future users would then be able to log on and navigate through the world, and the paths of their movement could be emailed back to the architect or researcher. In this way, the building model serves a dual purpose. Firstly, it acts as a method of liasing with and informing the would-be end users, a kind of public consultation exercise that is particularly approved of at the moment by Planning Authorities. Secondly, it becomes a means for the architect to generate data about how the real building might be used prior to its completion hence allowing time for design modifications before being constructed.

As it is essential to continue to determine whether or not there are similarities between real and virtual world navigation, then it is necessary to improve upon the current methods of real-world movement observations. At the moment, the kinds of data available from virtual experiments are far more detailed and precise when compared to real world observations that are mostly made by hand. For the kind of work undertaken in this thesis to continue, the quality of the real world observations must improve so that the two data sets are equivalent. The types of methods that can be used and developed for this purpose are techniques such as the use of active badges, video capture/analysis, infrared beams and pressure pads. This is a prime area for future research.

In all of the worlds that were used in this thesis' experiments, the subjects were walking around alone. The potential for using virtual reality techniques to investigate the effect of co-presence upon behaviour, particularly navigation, is a potentially exciting future direction. There are a number of ways of investigating this. Firstly, it could be possible to set up a series of observations inside online, collaborative worlds on the Internet (for example, AlphaWorld). They could be used to investigate the role of the environment and spatial design of the worlds upon navigation, chance encounter and casual meetings. This could serve to progress work already being conducted into this area by Huxor (Huxor 1998). Another direction that could be taken is to 'stage' a number of experiments. For example, if the following hypothesis were to be tested, namely that 'people simply follow people'. The testing of this could be achieved through the use of a world where certain streets were highly populated (with avatars) and other streets quite empty (whilst keeping all other variables constant) and then to determine if people choose to take the seemingly more 'popular' route over the choice of the less 'popular' route. This would effectively be testing the idea that people create a 'multiplier effect'; that we are more inclined to take a route if people are already present. It would also be interesting to determine whether there might exist gender differences in the results of such an experiment and it could be a useful way in which to investigate individual perceptions of personal safety.

It is hoped that in the next few years more sophisticated eye tracking equipment might be easily available. If it were possible to combine eye tracking equipment with a head mounted display then it would be possible to gauge precisely where subjects look as they are moving around instead of their approximate direction of gaze. This option was investigated at the beginning of the experimental work undertaken as part of this thesis, but it proved impossible to achieve. Technologically it is currently *just about* possible to make these kinds of measurements in both the real world and the virtual world. Such additional information could be particularly informative. For example, it could be possible to identify locations where people make route choice decisions and to determine what environmental features they look at in order to inform their decision-making. This information could not only be useful for urban designers but would have an immediate application for producing effective signage in large, complex buildings such as airports.

All of the above suggestions for future research were concerned with alternative methods of gathering real and virtual observation data. There are also a number of possible areas of research that arise from the methods of analysing such data. All of these future directions were prompted by work done in this thesis.

If much larger data sets were to be analysed, then it could prove useful to further develop the methods of cluster analysis used in this thesis. Ultimately, it could be possible to combine a variation of cluster analysis with methods of spatial analysis. For example, as it was suggested earlier in this chapter, it could be possible to test whether a pause point should be assigned to a cluster depending on whether it were visible to all other points already belonging to that the cluster. This would effectively combine a variant of cluster analysis with a measure of intervisibility/isovist relationships.

It has already been mentioned in this chapter that it could be possible to use virtual reality technology to investigate gender and age differences of navigational strategy. Another interesting area for research would be to determine whether or not there are differences between subjects with varying experiences of virtual environments. For example, to investigate differences between people who are particularly used to navigating through virtual worlds, such as computer-game players, and those for whom the technology is quite unfamiliar. This research particularly prompts the question of whether future generations who are growing up, utterly familiar with such technology, would ultimately demonstrate different behavioural characteristics of virtual navigation. Could these differences effect real world behaviours too?

One set of experiments that could arise from this thesis would be to investigate the effect of speed upon stopping behaviour. It would be extremely useful for designers of virtual environments to be able to determine if people pause for briefer periods of time if they are moving at a faster rate. This would be relatively straightforward to test. It could also be interesting to use virtual simulations of urban districts and buildings to determine their different effects upon micro-scale behaviours. Do we use urban space in the same manner that we use buildings, or does there exist identifiable differences?

Another pair of connected research areas is concerned with the generation of efficient representations of environments. Firstly, the potential for an automated environment-mapping tool is high, and for this reason, this should be regarded as a valid area of research stemming from this thesis. An automated environment-mapping tool could be based upon isovist analysis of environments, and then using this data as a basis, to identify and record the location of road junctions. Lines of sight and/or streets connecting junctions could be represented in graph-form or by the use of mathematical equations to describe junction-shape. The second area of work involves the further development of mathematical equations to describe road junctions.

A particularly interesting area of future research would be to investigate the phenomenon termed route angularity by using virtual worlds. In particular it could be possible to determine whether or not ‘seven plus or minus two’ junctions or changes of direction were significant. The use of virtual environments lends itself particularly well to this task, as it would be possible to generate a number of theoretical routes, all of precisely the same length, yet expressing a range of changes of direction. Once again, this would be a particularly straightforward set of experiments to create.

The analysis performed in Chapter 5, using string matching, could be further developed. In Chapter 5, the strings were used to represent the spaces through which subjects passed. It was also suggested in Appendix A that this technique could be applied to local decisions made by people, for example to represent the act of turning left or right. This could be used as a method for analysing the kinds of decisions being made. If string matching analysis (Chapter 5) were combined directly with route choice analysis (as in Chapter 6) and perhaps also combined with canonicalisation of angles (Appendix B) then the local strategy of people could be analysed as well as their global strategy. In this manner, someone who consistently took a left turn followed by a right turn could be regarded as being identical (in terms of local strategy) to a second person making the same turns but in a completely different part of an environment.

It could be interesting to perform the same kinds of experiments as performed in this thesis but to record other human-response data at the same time. Data that could be simultaneously recorded, and could broaden our understanding of how we navigate through environments, could be body temperature, heart rate, a subject's verbal discussions of their experience and the activity of certain parts of the brain. This is an approach that is already being taken by O'Keefe et al. on hippocampal activity during navigation (O'Keefe and Burgess 1996). There is, however, scope for more work to be done in this area.

Since this thesis investigated the micro-scale behaviours that result in the aggregate patterns of pedestrian movement that have been observed in buildings and cities, it could be possible to translate these micro-scale behaviours into a set of rules. These rules could be used as the basis to programme cellular automata or virtual 'people' (agents) who would be able navigate environments in a manner similar to real people. There are a couple of applications for this approach. Firstly, it could be a useful method to further pedestrian movement research. If the micro-scale behaviours recorded in this thesis could be translated into a set of rules and then the resultant patterns of movement were found to correlate with real pedestrian movement then this would validate the rule-set. The second application for which this could be used would be to populate virtual environments with virtual people who would be able to navigate in a natural manner. A criticism that is often levelled at architectural drawings and photographs is that they are often devoid of people. If these micro-behaviours were to be successfully translated into rules governing avatar behaviour, then architects could present a client images, movies and real-time walkthroughs of a design that was fully populated with virtual people. This could give the client a far better idea of how the final building would appear.

One of the more interesting measures to have emerged from the isovist analysis was 'drift'. Drift is the distance in length of a line connecting the viewpoint of an isovist to its centre of gravity. The patterns formed by areas of minimum drift tended towards the centres of spaces and roads were incredibly suggestive of patterns of pedestrian movement, sufficient paths and axial lines. It would be particularly interesting to investigate whether drift could be connected to sufficient paths and axial analysis.

Finally, the methods introduced in this chapter for analysing environments using isovist arrays, and correlating this data with micro-scale behaviours of people could, without question, be continued and developed. When Benedikt discussed the potential for conducting research of this kind in his paper (Benedikt 1979), it was particularly difficult to record adequate data of people's behavioural patterns in such a manner that they could be correlated to isovist properties. The use of virtual environments to achieve this goal has been suitably demonstrated in this thesis and it is to be hoped that this work may be continued.