Chapter Seven:
Pause Point Data and Direction of Gaze Analysis

Abstract

This chapter will illustrate how the same raw data used to determine subject paths could also be used to calculate 'pause points'. A definition of a pause point is that it is a key location in the world where people stop or pause for a significant amount of time. The patterns formed by such pause points can be generated for different time durations. After an initial explanation of how pause points are calculated, the pause points for the seven test environments are calculated. After presenting the results of the pause points for all the worlds, a method for identifying clusters of pause points using the k-means cluster analysis is introduced. The results of the ensuing cluster analyses are then presented in full, for each world, with appropriate interpretations. Individual differences (between subjects) in patterns of pausing is explored by considering the numbers, frequency and total duration of all pauses for one of the worlds. The second half of this chapter investigates the approximate direction of gaze of the subjects while they are traversing through an environment. The first and less satisfactory method of direction of gaze analysis records the continuous direction of gaze of a subject throughout their journey. This method is superseded by a second method focusing only on the direction of gaze of a subject whilst stationary. This method of analysis effectively combines both the pause points and direction of gaze representations. Therefore, the question that emerges from this chapter is, where precisely are people pausing and what are they looking at whilst stationary? The direction of gaze analysis covered in this chapter leads directly to methods of analysing the visual fields (isovists) of these virtual worlds and their correlation to pause points. This is the topic covered in the following chapter, Chapter 8.

Introduction

In Chapter 3 and at the end of Chapter 4, the raw output data (ASCII log file) from the virtual experiments were represented as simple paths taken by individuals through the environments. This is a representation of the behaviour of a subject (or group of subjects) at its crudest and most basic level. The diagrams at the end of Chapter 4 purely illustrated *where* people had travelled. Earlier in this thesis it was suggested that this could be regarded as a 'top down' approach to answering the question of which small scale actions might be important in the production of observed patterns of real-world movement. This chapter begins to make a shift from examining continuous routes and begins to investigate what individual actions may be yielded by the same data. Tactically this can be held to be a 'bottom up' approach to the same question; the focus of this chapter is upon smaller-scale actions not routes.

In Chapter 3, as part of an explanation of how the first routes were visualised, a list was presented which suggested the type of additional information regarding a subject's journey that it could be possible to extract from such data. This information was as follows:

- Continuous path of movement of any person
- Location of a person at any moment of time
- Location of a person's 'Pause Points' or 'Dwell Points'
- Number of a subject's 'Pause Points' or 'Dwell Points'
- Proportion of the journey spent stationary by any person
- Total distance travelled & total journey time
- Average velocity/acceleration over the journey
- Continuous locus of the orientation of the head mounted display
- Approximate direction of a subject's gaze whilst stationary

The first two items on this list have effectively been covered in the proceeding chapters. This chapter will illustrate how it is possible to calculate and represent the majority of the measures listed above. This is an example of one technique for moving from the macro level (the movement paths) to the micro level of smaller scale behaviours that take place en route.

The Calculation of Pause Points

In considering individual actions such as pausing behaviour, the questions that are prompted are where precisely along a journey does a subject pause, at what kinds of locations does a person stop and for how long? In order to calculate pause points from the original log file data it was necessary write a small computer programme[1]. This programme was based upon the earlier application that transformed the raw ASCII log file into a 2d CAD representation of the paths. However, it was necessary to adapt and extend the original programme to calculate where people were pausing in the environment and to represent these locations as a 'point' on plan. These points are the graphical indication of pause points or dwell points.

How are these dwell points calculated from the log file? The raw data are simply series of points in space that are occupied by a subject and ordered chronologically. The first stage is to compare each sampled position to its neighbouring locations (in the list) and to calculate the distance moved by the subject between every two consecutive sampled points. If that distance is zero, then the subject has been stationary between those two sample-points. By comparing each world location in the list to the preceding and following positions to record any locations of zero movement, a separate sub-list of dwell points can be compiled. This is only the first stage. It then becomes necessary to define what is actually a *deliberate* pause and what might be simply a momentary hesitation, caused, for example, by a finger slipping off the 3D mouse. If a person pauses for a fraction of a second, should this be counted as a dwell point? The next stage, therefore, is to establish a *lower limit* below which a person is not held to be deliberately pausing. There is no need to establish an upper limit; a person will simply remain at the same location until such time they decide to move on. It is the lower-limit time 'trigger' that is the most important criterion to establish.

A decision must be made to determine which time 'trigger' is to be used, such that any pause *longer* than this amount be counted as a deliberate pause. The strategy used to determine which time triggers were significant was to initially calculate ten-second, five-second, two-second and one-second pauses for all worlds. Note that, in this thesis, the phrase 'x-second pause' indicates that the total duration of a pause point can be any length of time longer than and including x seconds. The results of calculating a range of pause points for all worlds were that there were almost no occurrences of ten-second pauses and relatively few five-second pauses (other than at the start and finish locations). In contrast, in most worlds there appeared to be an unwieldily high number of one-second pause points. For most worlds, the number of two-second pause points fell approximately midway between the totals of one and five-second pauses and intuitively it seemed that this time interval was the optimal one to use for the majority of the environments. Unless stated elsewhere, this is the pause point duration being calculated. It may seem as if this is a very brief interval of time, but it is suggested that conceivably such duration feels longer in a virtual world than in the real world. It could be that a two-second pause might be analogous to a much longer pause in a real environment and this could be a useful area for future research.

It should also be stressed once again that this time duration is, in fact, the *minimum* time interval recorded and that many subjects were in fact stopping for periods far exceeding this amount. Figure 7.1 below shows a typical world, World C and the path of a single person superimposed by their two-second pause points, indicated as points along their route.

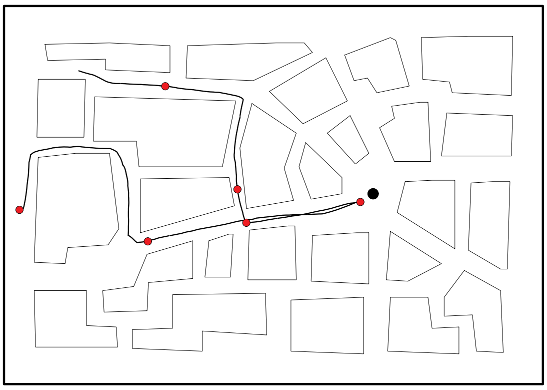

Figure 7.1 A Single Route with Pause Points

After calculating and representing the patterns of pause points for *one* subject in an environment, the aggregate patterns of all pause points for all subjects in each world can be produced. This was achieved by superimposing all pause points for all subjects on the same plan. On calculating the aggregate patterns of pause points for all seven worlds, it was eventually decided to use different durations of pause points for some of the environments (although two-seconds is still the predominate time-duration used). Unless stated otherwise, the pause point duration is two-second. However, whenever a decision was made to use alternative pause point durations, full explanations for the reason are given alongside the interpretation of the results (in the following section). A summary of the pause-point durations for the different worlds is given in table 7.1 below.

World	Pause Point Duration (minimum limit)
World A	5 second
World B	2 second
World C	2 second
World D	2 second
World E	0.5 second
World F	0.5 second
World G	0.5 second

Table 7.1 Pause Point Durations for all Worlds

Patterns of Pause Points in the Experiment Worlds

The following set of figures shows the patterns of pause points for all subjects in all seven, test environments. An explanation of the results are given, discussing each world in turn.

World A, The Tate Gallery

In World A or the Tate Gallery, the pause points illustrated above are calculated at five-second instead of two-second intervals. There were a far greater number of two-second pause-points in this world compared to the other six worlds. It was for this reason a longer time interval was used. It would appear that we may move at a different pace in a building

$$\left(\frac{Total\ Duration\ of\ all\ Pause\ Points}{Total\ Duration\ of\ Whole\ Journey}\right) = Propo$$

Figure 7.2 Pause Points in the Tate Gallery

than in an urban environment, even through these environments are only simulated. It would be logical that the difference in scale between buildings and urban environments could produce differences in patterns of stopping behaviour. A time interval that constitutes a significant pause at the building level may not constitute a significant pause at the urban level.

It is initially conjectured in this chapter that people could be pausing at locations where route choice decisions need to be made. At the urban level, these locations are obviously road junctions. At the building level analogous locations could be the thresholds between adjacent spaces, such as doorways. In which case it could be expected that subjects should be pausing in or in close proximity to doorways. Although some pause point locations near doorways can be observed, this does not seem to be the overall pattern of locations. Quite often, a subject appears to be stopping in the centre of the room and looking around. However, this behaviour may be attributed to the scale of the environment. At the urban level, where junctions are often situated at a considerable distance from other junctions, it is not unusual to be unable to see around a corner until quite close to a junction. In complete contrast to the urban experience, due to the smaller scale of a building, it is possible to stand in the centre of a space and on scanning the room, to catch penetrating glimpses through a *number* of doorways simultaneously, namely that in a building it is possible to survey multiple route options simultaneously in a manner quite impossible at an urban level (assuming that doors are open).

It can also be noted that there are a large number of pause points in close proximity to the subjects' starting position, at the main entrance (the bottom, middle on plan). The density of pause points remains high in the vestibule beyond the entrance lobby and only begins to fall off after this second entry space.

World B

In World B, the subjects all started their journeys in the uppermost, left-hand corner, on plan. The exact location of their starting position was in the centre of the open 'square' formed between the corner of the boundary wall and the concave corner of the adjacent cruciform block and there are a large number of pause points situated around this area. It soon became clear that it was usual for a subject to remain stationary upon first entering a world, whilst they looked around in order to both acquaint themselves with and orient themselves to the environment. Another obvious location for pause points is the centre of the world, which is, in this particular case, also the wayfinding goal. Again, it appeared to be quite common for subjects to pause to scrutinise the monument upon reaching their goal, perhaps to confirm that they had really reached their destination. Apart from these two patterns of stopping, the only characteristic shared by the pause point locations in the other worlds is that they appear to be in close proximity to road junctions. The relative density of pause points is greatest along the primary route from the starting location to the wayfinding goal and in direct proportion to the distance from the starting location, namely that early on in a subject's journey they appear to be more inclined to pause momentarily (perhaps whilst familiarising themselves with the controls).

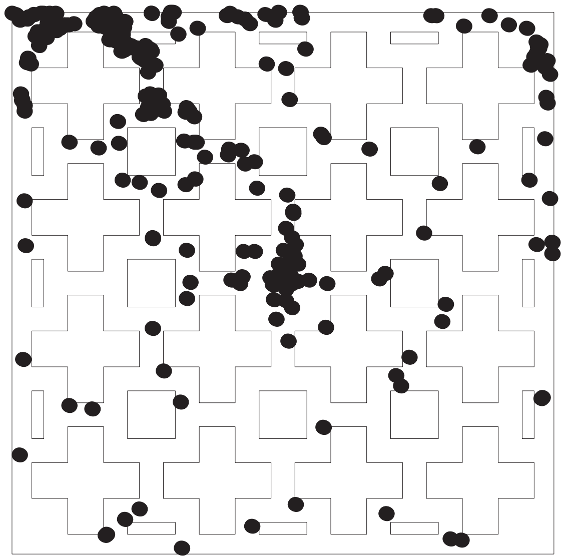

Figure 7.3 Pause Points in World B

World C

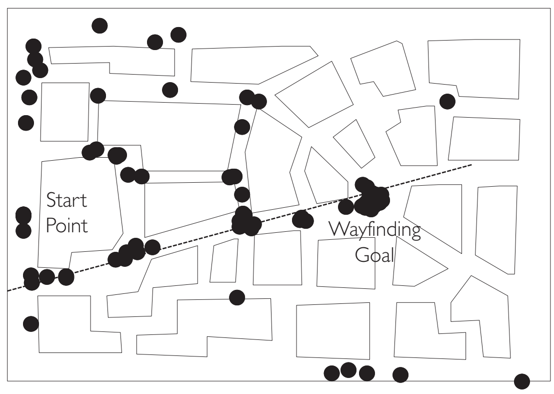

Figure 7.4 Pause Points in World C

In World C, the subjects start their journeys at the middle, leftmost side of the plan view. This location is marked by a number of pause points, which appear as only one or two points, since they are coincident. A large grouping of pause points is evident around the monument, in the central 'square'. This was the wayfinding goal used by the subjects. Again, there is a greater number of pause points scattered along the most direct route from the starting position to the monument, namely the straight, diagonal route to the lower half of the plan view (indicated with a dotted line). Again, the locations of the pause points appear to be distributed in close proximity to road junctions.

World D

In World E, the subjects started their journeys in the same location as the previous world, World C. As in World C, there are also a large number of pause points located around the monument. There appears to be fewer pause points in this world, spatially distributed more evenly throughout the environment. Again, it does appear that subjects are pausing at or close to road junctions. At first, it appears puzzling that people are pausing less in the more unintelligible and potentially more confusing world. It could be expected that subjects would pause more frequently, the more disorientated or lost they become. Instead, the opposite appears to be happening, as subjects become disorientated they pause less. It may be that subjects are pausing only in locations that afford them maximum environmental information. If a subject is lost in a segregated, visually limited part of an environment, they do not pause to look around. Instead, they continue until they emerge into a more integrated, larger space, at which point they then pause to take stock and re-evaluate their route-choice decisions. It then follows that subjects would stop less not more in an environment that was intrinsically confusing (i.e. unintelligible).

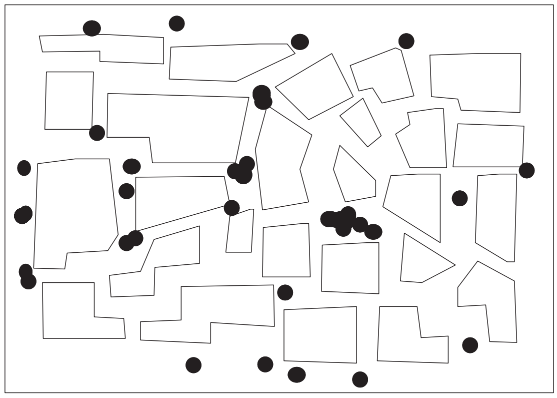

Figure 7.5 Pause Points in World D

World E

Worlds E, F and G were the last experiments to be performed. Due to problems with the equipment set-up, by this stage, many of the subjects were feeling slightly nauseous. In order to enable these subjects to complete the experiment, all subjects were shown how to accelerate their pace, enabling them to move at a faster pace. (See Chapter 4 for a fuller explanation of the experiment procedure.)

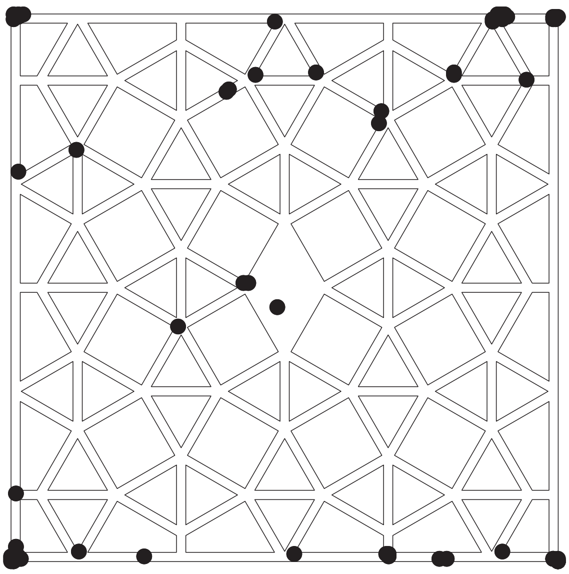

Figure 7.6 Pause Points in World E

Surprisingly, the patterns of pause points for these worlds seem to reflect the increased pace. Since the subjects were moving at a faster speed, they appeared to pause less often and for smaller time intervals. On subsequent calculation, there were very few two-second-pause points and therefore it was decided to attempt to calculate half-second pause points instead. The resultant pattern of half-second pause points appears to follow the expected pattern for longer pause points, based upon the results of the other worlds. The patterns of stopping behaviour for worlds E, F and G use half-second pause points rather than two-second pause points for their calculations. This ensuing behaviour had not been anticipated during the experiments, the increased pace having only been introduced as an impromptu measure to ensure the completion of the experiments by all subjects. However, both Worlds E and F appear to confirm the observation that subjects appear to be pausing at locations where a route choice decision needed to be made. It also appears that this pattern is more marked in these two worlds than in any of the other worlds. See World F results in figure 7.7 below.

World F

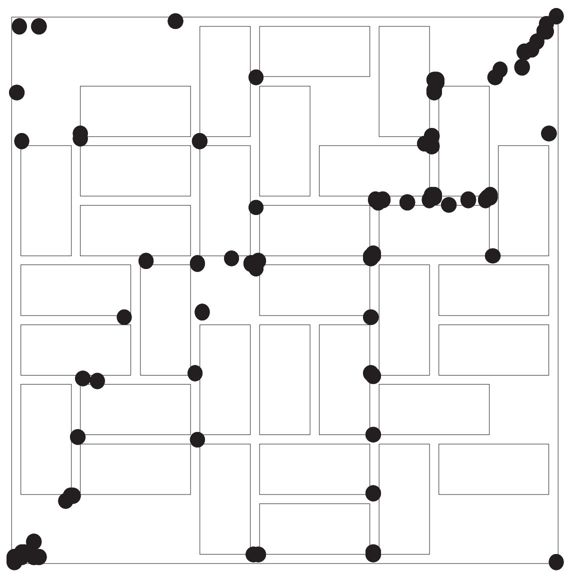

Figure 7.7 Pause Points in World F

World G, Barnsbury

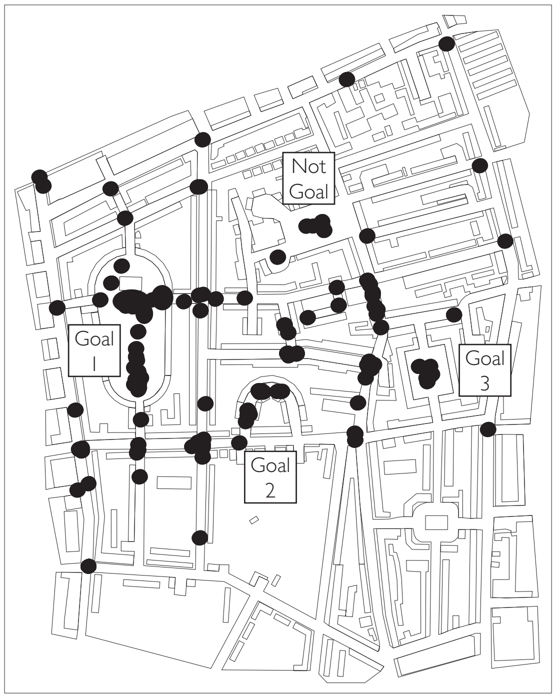

Figure 7.8 Pause Points in World G

In Barnsbury, all subjects began their journey in the bottom-leftmost corner, marked by what appears to be a single pause point (in reality a large number of coincident pause points). There are clearly a number of pause points clustered around the three wayfinding goals, goal 1, 2 and 3 on the plan (the keyhole shaped ‘square’, the crescent and the square). Another square (also marked on the plan, in the upper, rightmost quadrant) proved to be quite a popular place to stop. Many subjects paused here to determine if they had found goal 3. If points lying outside the areas marked as either goals or potential goals are considered, it can be seen that most of these locations also appear to be located in or close to road junctions. This is entirely in keeping with the results found in all of the other experiments.

All these worlds appear to confirm that people are pausing, in the main, at locations where a route choice decision needs to be made. This would be quite a 'common-sense' finding to present. It would be logical to conclude that this should be the behavioural pattern of groups of subjects within an unfamiliar environment. However, this is purely an initial interpretation based upon a visual presentation of the subjects' pause point locations. In order to make a more accurate assessment of the patterns of people's stopping behaviours, a more accurate method of determining these patterns must be devised. Since it would appear that people *might* be responding to visual cues, then a method of interpreting this pause point data could be developed using isovist analysis. This is covered in the following chapter, Chapter 8.

Pause Point Cluster Analysis

In the previous section, it appeared that there were identifiable groups of pause points. In most worlds, subjects appeared to be pausing at the start of their routes and upon reaching their wayfinding goals. It was also conjectured that subjects appeared to be pausing near or in close proximity to road junctions. Alternatively, as in World A, it was suggested that subjects appeared to be pausing at the centres of rooms. Since this thesis is concerned with which small-scale actions might be important, then the identification of *any* consistent patterns in subjects' pausing behaviour could contribute to an understanding of how small-scale actions produce the emergent phenomenon which is pedestrian movement. The question is, do clusters of pause points occur more in some regions of the world than in others?

In order to determine whether or not there might be any clustering pattern to the location of the pause points, it was decided to perform a form of cluster analysis on each aggregate set of pause points. The method used is a relatively simple method of cluster analysis termed k-means clustering; k-means is a straightforward and efficient process of establishing clusters using an iterative algorithm.

The K-means Clustering Algorithm

This algorithm was first developed in 1967 by MacQueen. The k in k-means indicates the number of clusters to be found. This is usually established before any calculations are made. The k-means algorithm is an iterative algorithm that means that it will run repeatedly until a solution is found, generating more accurate approximations to the correct answer each time.

Imagine a number of data points. Let the total number be n. The number of clusters, k, is first established (i.e. k must be given a value before the algorithm can run) and it is essential that $k \leq n$.
A number of data points are chosen randomly from the sample, this number being equal to the number of clusters being sought. For example if there are one hundred points in the data set and an attempt is being made to seek five clusters, then five points are chosen at random from the original one hundred. Each one of the randomly chosen points acts as the centre of a 'proto-cluster'. The centre of a cluster is the mean value of all the points or,

$$Centre(x, y) = \left[\sum\left(x^1, x^2, x^3 \ldots x^m\right)/m\right], \left[\sum\left(y^1, y^2, y^3 \ldots y^m\right)/m\right]$$

Equation 7.1

where m is the number of points in a given cluster.

This can be held to be analogous to 'seeding' the clusters. A test is then performed. The Euclidean distance between each point in the data set and each of the k number of cluster centres (the randomly chosen points) is calculated. When the centre closest to a point is identified, it results in the point becoming a member of that cluster. This process is continued until each of the data points is a member of one of the k number of clusters. All points must be assigned to a cluster. When this stage has been reached a new centre for each cluster can be calculated based upon the mean location of all the data points forming each cluster. A new set of cluster centres is then calculated (as opposed to being randomly assigned as in the first example). The whole process of determining which is the closest cluster centre is then repeated, with the data points being re-assigned. Some points might become members of alternative clusters.

Each time this process is repeated, the new cluster centres are calculated (and new clusters formed). However, each time a new centre should be slightly closer to its old centre (and this distance should keep diminishing with each iteration), until finally there is no change in location between one cluster centre and its previous or subsequent centre-locations. When the cluster centres have settled down to a stable state the k-means algorithm has been completed. This is why it is termed a 'self-organising' algorithm.

The only difference between the implementation of the algorithm described above and the one used to produce the following clustering diagrams, is with regard to the assignment of k, the number of clusters. As stated above the first stage in the k-means calculation is usually to decide how many clusters are being formed. In the statistical package, Datadesk, used for these examples it calculates all possible numbers of clusters from 1 to n. It is possible for a point to be both a member of a cluster of 1 (it is a cluster on its own) and hence there are n clusters, or for a point to be a member of one giant cluster with all other points as co-members and no other clusters. All other possibilities in between are also calculated. The results are presented as a tree diagram.

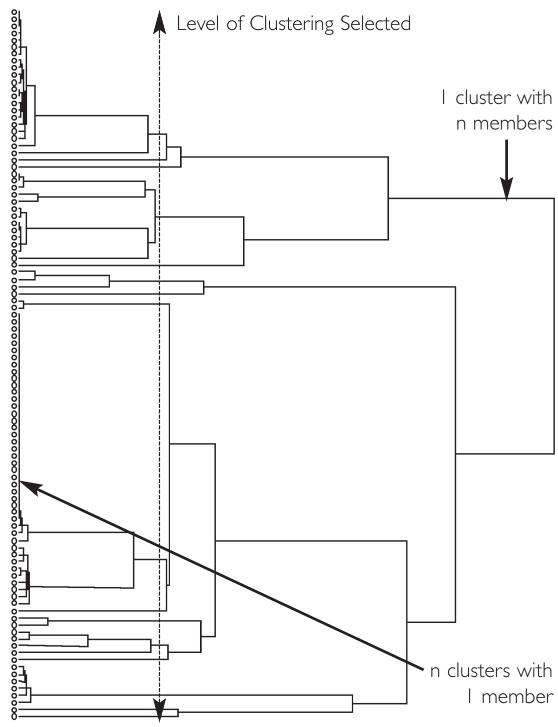

Figure 7.9 Cluster Tree Diagram

By using the tree diagram above, a level of clustering was selected (by eye), that was neither too extreme (in either direction). This stage is indicated by the pointed section line marked on the figure above.

The results of the k-means clustering analysis, performed on each of the seven virtual worlds are presented overleaf.

World A, The Tate Gallery

Figure 7.10 Pause Point Clusters in World A

World B

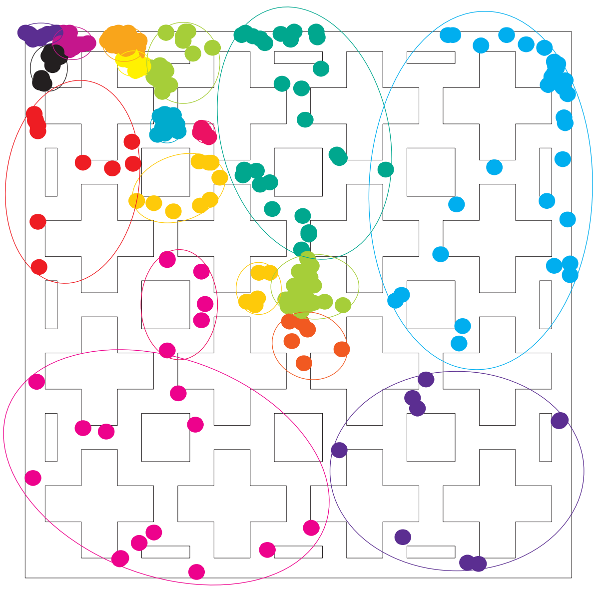

Figure 7.11 Pause Point Clusters in World B

World C

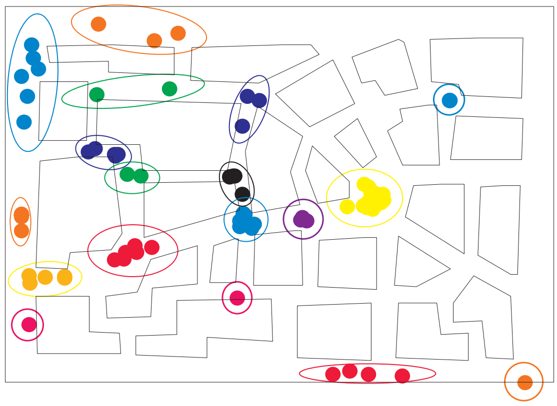

Figure 7.12 Pause Point Clusters in World C

World D

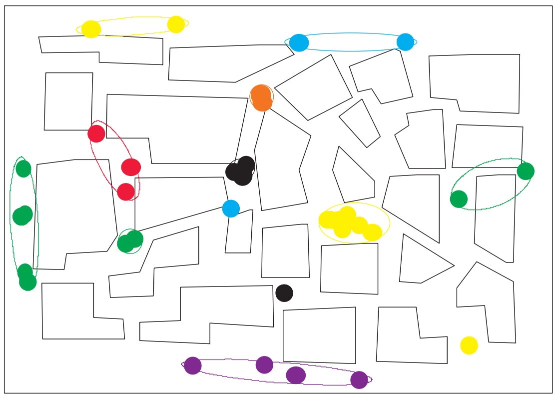

Figure 7.13 Pause Point Clusters in World D

World E

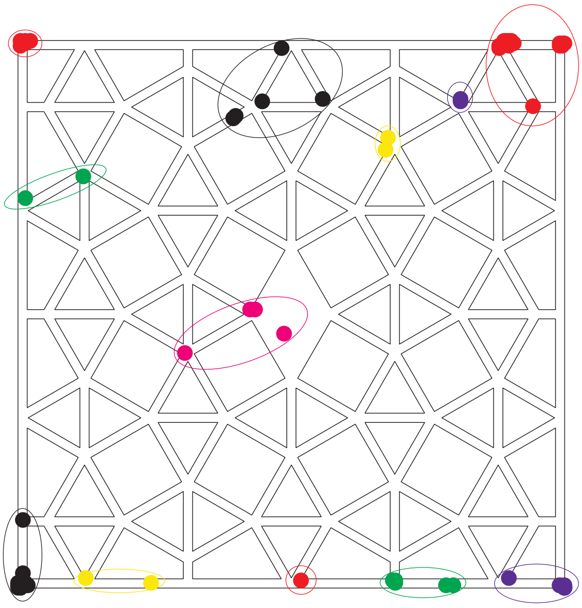

Figure 7.14 Pause Point Clusters in World E

World G

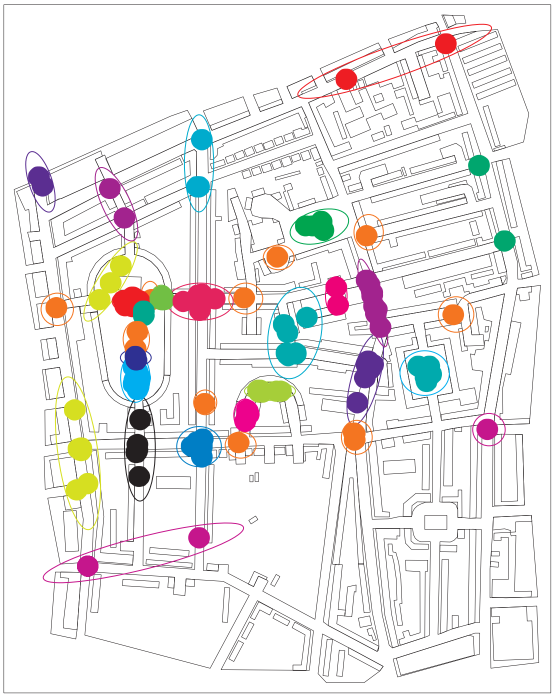

Figure 7.16 Pause Point Clusters in World G

World F

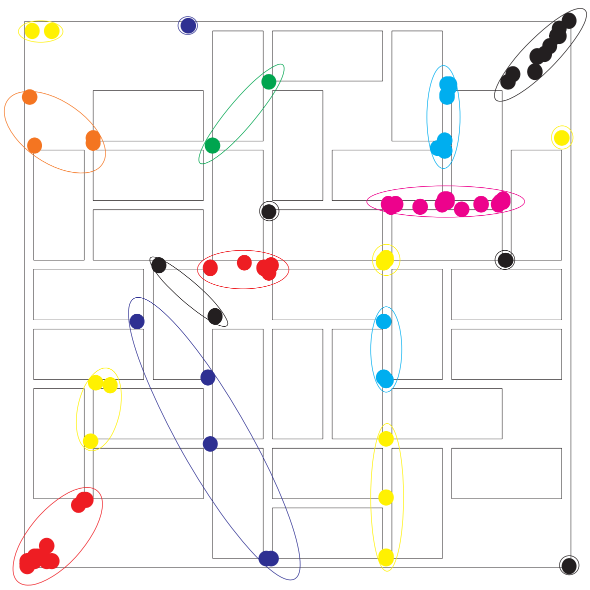

Figure 7.15 Pause Point Clusters in World F

There is no need to provide additional descriptions of the clustering diagrams above, since they are quite self-explanatory and any descriptions would only mirror those provided alongside the basic (non-clustered) pause point analysis in the previous section.

Considering all the pause point clustering diagrams above, the k-means algorithm appears to be quite effective at identifying clusters; it frequently identifies clusters of points, which intuitively one would have identified by eye. In this manner, it could be used to automate the analysis of pause point patterns and could be used to calculate the centre of a group or cluster of such points. However, in this particular data set, it does little to highlight any patterns not visible to the eye; it serves only to support one's intuition rather than augment it in any way. In many respects, it would appear that the value of such a method of analysis rests in the interpretation of the results. In this particular case, any interpretation would simply mirror the initial descriptions of the patterns of pause points given in the previous chapter section. This is not to say that using a type of statistical cluster analysis *only* mirrors our intuition. There are occasions when patterns not visible to the eye can be highlighted or identified by using this method. It is simply that in the examples shown above, this method of analysis seems to shed no new light onto the results.

Occasionally, in the above examples, the cluster analysis identifies a group of points which, considered against the urban street or building configuration, appears to be counter-intuitive. Since k-means cluster analysis *only* uses the co-ordinates of the pause points and does not take into account the urban/building structure in any manner, it is questionable how useful a method of analysis it is. The algorithm will sometimes identify a cluster of points, which in reality is quite meaningless. Clearly, when taken in context, certain pause points may have nothing in common whatsoever. This flaw could be easily rectified by adapting the k-means clustering algorithm to include spatial information. Currently a point is assigned to a cluster based purely upon its distance from that cluster's centre. However, a second condition for joining a cluster could be introduced. A point may join its nearest cluster if the centre of the cluster is visible from that point. If not, it is assigned to the next nearest cluster (providing its centre is visible) and so on. If no clusters are visible to a point, it becomes a new proto-cluster in its own right. In this way, although a number of clusters would need to be pre-defined before running the algorithm, it would be impossible to predict how many clusters would actually be formed once the algorithm had reached a steady state.

In conclusion, if the string matching algorithm presented in chapter 5 is a way of looking for patterns or similarities amongst a sample of routes, then cluster analysis can be regarded as performing an analogous task for pause points. Although each method is presented in a rudimentary manner in this thesis, there is a need to develop a set of more sophisticated tools to enable this kind of pattern recognition. If it is to be possible to calculate how aggregate pedestrian movement arises from a number of small-scale actions. It is not only necessary to be able to observe and identify these actions, but more importantly to seek patterns and similarities amongst them. The work presented in this chapter and chapter 5 can be regarded as first steps towards this goal.

Numbers of Pause Points, Travel Distance & Proportion of Time Spent Stationary

Other patterns might be found in the distribution of pause points. Instead of examining how the pause points are *spatially* distributed in an environment, it may be possible to discover other relationships in the number or frequency of pause points over a single journey. This could lead to other questions, such as, can similarities be found between various groups of individuals? Therefore, the next set of information to be extracted from the log files, is the total number of an individual's pause points. These can be correlated to the proportion of a journey that is spent stationary by any person and to either the distance travelled by a person or their total journey time. A single example of this type of data analysis is worked through to illustrate how it may be analysed in this way.

If we use World B as the example, it is possible to take the log file and process it in order to extract the number of pause points for each subject. One hypothesis raised by this form of analysis is that in an unfamiliar environment, where a subject does not have prior knowledge of their route, they might be likely to pause more frequently during their journey. In this world, the least number of times that a subject paused was twice, the greatest number was twenty and the mean number was 9.4 times.

The development of methods to analyse these dwell points may potentially provide additional information on how a subject navigates through an unfamiliar environment. In particular, analyses of both the location and duration of these dwell points might illuminate which aspects of an environment the subjects uses to aid such navigation.

Duration of Dwell Points

To examine the amount time spent stationary by each subject in World B, this method selects the *proportion* of time spent stationary for analysis, in preference to the total amount of time spent stationary (measured in numbers of points sampled[2]). In a long journey, it would be expected that the amount of time spent stationary would be greater, since the entire journey time itself is greater. The proportion of time spent stationary, is independent of the journey time, since it considers the length of journey. The proportion of time spent stationary, can be calculated as follows.

$$\left(\frac{Total\ Duration\ of\ all\ Pause\ Points}{Total\ Duration\ of\ Whole\ Journey}\right) =$$

$$Proportion\ of\ Journey\ Spent\ Stationary$$

Equation 7.2

This method also has the advantage that by using proportions the result is dimensionless. It does not matter if the calculations are done using numbers of samples or seconds, since the result is identical

In the sample of 36 people undertaking the I.C.A. experiment, the minimum amount of time spent stationary by any subject was 11% and the maximum amount of time spent stationary was 70%. The mean amount of time spend stationary by all subjects, as a proportion of the total journey length, is 26%. Ttable 7.2 overleaf contains the tabulated data of these calculations. Once the percentage proportion of time spent stationary has been calculated for each individual subject, this data may be plotted against the total journey length.

All subjects were given the same instructions, namely to locate the monument at the centre of the virtual world. Therefore, it can be assumed that those subjects who found the monument very easily were those who took the most efficient (most direct) route. It can also be assumed that these subjects would also have the least journey distance, as measured in metres. Conversely, this implies that subjects who became the most lost and encountered the greatest difficulties in finding the monument were the ones who took the least efficient routes and had the greatest journey distance.

The comparison of the percentage of journey time spent stationary to the total journey length measured in metres can be used to test a hypothesis. The hypothesis being tested is that the more lost a subject becomes they more they pause to try to work out the correct route. If this hypothesis were correct, then there would be a positive correlation between the length of journey time and the proportion of pause points. In fact, if we examine the data from

Subject Code Number	Journey Duration / Number of Samples	Duration Spent Stationary / Samples	Total Number of Stationary Locations	Distance Traveled / Metres	Proportion of Journey Stationary
1	729	131	8	409	0.18
2	868	208	7	112	0.24
3	836	202	15	439	0.24
4	1086	194	9	647	0.18
5	536	165	5	218	0.31
6	677	337	6	227	0.5
7	570	84	6	350	0.15
8	558	292	12	181	0.52
9	1214	401	18	512	0.33
10	780	260	20	316	0.33
11	486	124	10	234	0.26
12	1180	181	10	682	0.15
13	446	180	8	179	0.4
14	437	130	11	190	0.3
15	1030	719	4	210	0.7
16	1276	137	10	799	0.11
17	1141	276	19	574	0.24
18	1273	188	12	785	0.15
19	431	131	6	202	0.3
20	750	104	9	291	0.14
21	797	196	12	399	0.25
22	436	48	2	267	0.11
23	681	75	5	329	0.11
24	481	153	10	195	0.32
25	425	102	3	209	0.24
26	461	99	8	212	0.21
27	1057	121	13	659	0.11
28	664	347	19	201	0.52
29	443	85	4	246	0.19
30	355	93	5	179	0.26
31	505	58	3	275	0.11
32	596	81	7	358	0.14
33	692	263	19	254	0.38
34	787	129	5	459	0.16
35	917	243	12	410	0.26
36	440	73	6	251	0.17

Table 7.2. Tabulated Dwell Point Data

World B we find the complete opposite. Rather than a positive correlation between the two, we find a negative correlation, namely that those subjects who became more lost, were those who paused least to find their way. In the following section the data and scattergrams will be examined in detail in order to discuss why this may be so.

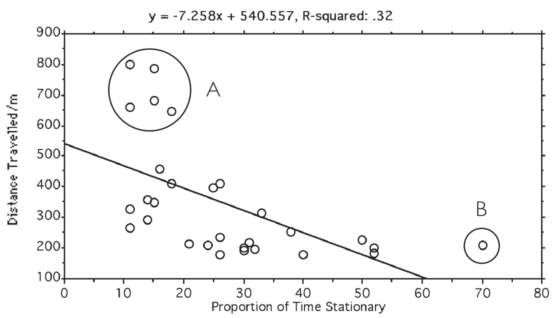

Figure 7.17 Scattergram of the Relationship between the Proportion of a Journey spent Stationary and Distance Travelled

If we consider first the five points circled on the upper top left hand corner of the scattergram, we find the mean proportion of time spent stationary by these five individuals is 14%, comprising of measures of 11%, 11%, 15%, 15% and 18% respectively, i.e. these subjects had some of the shortest dwell point durations of all of the subjects in the sample (11% is the least proportion spent stationary of all subjects). These same five subjects[3] also took a significantly longer time to reach their goal, measured in distance travelled. There are no other subjects who took longer or more circuitous routes to the monument. It appears that those subjects who found their goal most efficiently, i.e. by traversing the least distance to get there, were also those who paused the most.

It was not found that subjects who became lost compensated for their disorientation by pausing more frequently to find their way. Rather, it can be conjectured that those subjects who had most difficulties finding their goal, did so *precisely* because they paused infrequently and hence potentially missed vital visual clues indicating the correct route. In fact, by inspecting the subjects' routes on plan, at least two instances can be found, where subjects passed within sight of the monument. They would have seen their goal if they had only paused to look around but instead they moved past, evidently missing their goal.

If we now return to the scattergram and examine another anomaly, there is a single point, circled on the bottom right hand side of the scattergram. Here we find a subject who spent a disproportionate amount of time stationary. In fact, this particular subject spent 70% of her time stationary. We can also examine the data and extract the number of locations this subject paused; the subject only paused in four locations. (The average is 9.4 times for the whole sample.) What personal information might explain the behaviour of this individual? The subject was female and was the oldest person to undertake the experiment; her age was 52 (the mean age of the sample was 33 years). In contrast to the other subjects, she also reported that she experienced feelings of dizziness and nausea and that she consequently had problems navigating. In conclusion, many factors about this subject's personal information and experiences of the experiment appear to be very different to the group of subjects as a whole. If we exclude this subject from the sample and plot the new scattergram, we find a slightly tighter correlation. See the f igure 7.17 below.

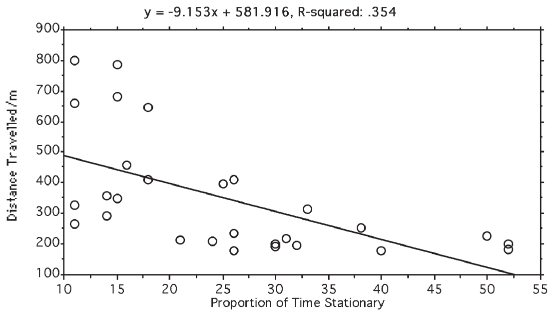

Figure 7.18 Scattergram of the Relationship between the Proportion of a Journey spent Stationary and Distance Travelled

In conclusion, it appears that there is a slight negative correlation (r2 of 0.32 or 0.354 excluding the subject who spent 70% of her time stationary) between the proportion of a subject's journey time spent stationary and the length of their journey measured in metres. Another finding of this method of data analysis is that this method can potentially be used to identify subgroups of subjects. This can be achieved by identifying similarities in subjects' spatial behaviour (i.e. all of them getting lost yet hardly pausing to consider their environment). It may be that with a much larger sample size and using different test environments, it may be possible to identify many more subgroups of people based on common spatial behaviour. This could be a valid area of future research, but is beyond the scope of this present thesis.

Direction of Gaze

Since it appears that subjects may be pausing at junctions, at locations where they have to make a route choice decision then an examination of their approximate direction of gaze might afford an insight into which visual cues are contributing to subjects' decision-making process.

Up until this stage in the chapter, all methods of analysis and representation have been calculated using the first three values in the log file, (see page 53, for an example of a log file), the x y and z location of the subject, at any instant of time. The following two techniques for investigating micro-behaviours en route, use the final three values. These three values represent the amount of rotation of the head-mounted display about the x y and z-axes. In other words, by an analysis of these angles, the orientation of the head set and hence a subject's head can be determined at any point of time. This can be held to be analogous to the approximate direction of gaze of the subject at any moment in time. Since the field of view of the head set is also known, which for this particular head set is 105°, then the precise field of view of the subject can also be recreated at any instant during their journey. The analysis of the varying field of view along a route is covered in the previous chapter, Chapter 6.

The primary question in this section of the chapter is to what extent are such head movements indicative of the actual direction of gaze or eye movement of a subject. In an important paper on predictable eye-head coordination during driving, Land (Land 1992) describes an experiment that he conducted upon subjects in a simulated driving environment. Whilst driving he measured both eye movements and head movements and found a high degree of congruity between the two.

Classical eye movement research typically involves the head being immobilized in some manner, such that it is only the movements of the eyes that are captured. Rarely have both head and eye movements been analyzed concurrently. The remarkable semblance between the head and eye movements found by Land, led him to conclude that both actions arose from a single 'direction of gaze' command. Other work in the eye movement field suggests that we rarely move our eyes more than 15° within their sockets[4] and that movement, outside this range, is instead accounted for by movement of the head. Based upon Land's work and in the absence of further research to the contrary, it could be reasonably suggested that the head movements of the subjects captured in the head set *are* indicative of the actual direction of gaze, given a small amount of error correction[5].

The first method of analysis presented in figure 7.19 overleaf analysed *all* head movements throughout a subject's journey. This data is represented by an arrow on plan, indicating the approximate direction of gaze of a subject. The locus of the direction of gaze is also indicated by a dotted line on plan.

This method of representation was felt to be less than satisfactory. This was because of the fact that although this method of visualising the data presented an accurate representation of the continuous head movement of a subject whilst navigating a virtual world, it was not particularly informative. The way in which Division's Divisor headset is configured, means that a subject navigates by looking in the direction in which they wish to travel and then using the 3d mouse to start moving. In this manner, whilst moving, the movement of the head and direction of gaze is tightly bound to the subject's method

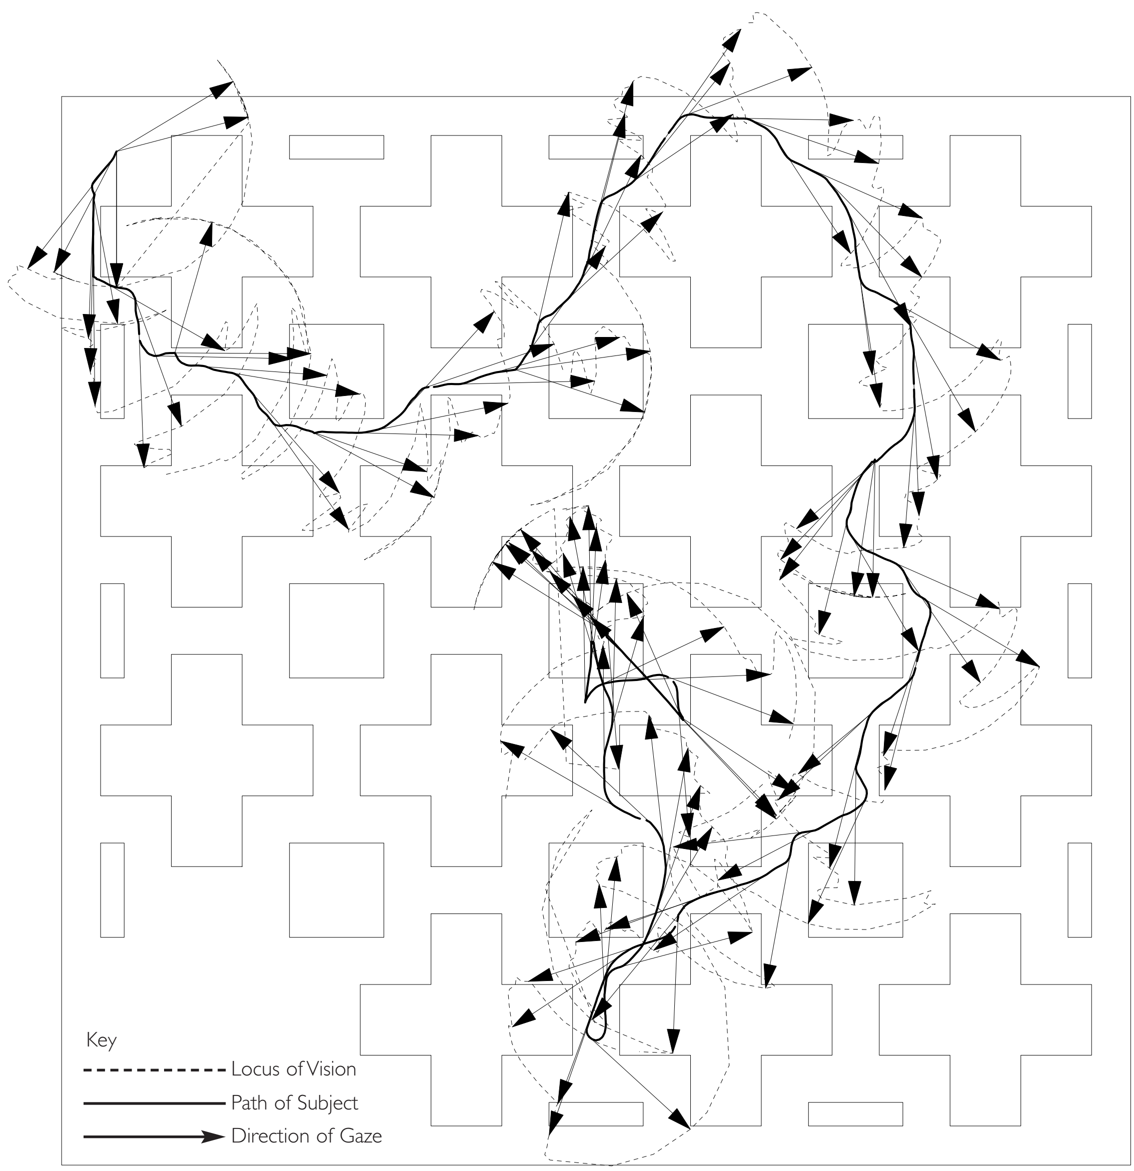

Figure 7.19 Path of a Subject through World A, showing Path, Direction of Gaze Arrows and Continuous Locus of Vision

of navigation. It was possible to reconfigure this arrangement, for example to allow a subject to move in the direction that they were pointing (with the 3d mouse). However, if the equipment were configured in this way, it would have been possible to move in one direction whilst looking in a completely different direction, even backwards. Trials conducted that experimented with severing the link between head orientation and direction of movement appeared to result in a less intuitive and harder to use interface. It was decided to continue to use the default configuration.

Instead of re-configuring the hardware a second method of analysis was developed which concentrated purely on the head movements whilst the subject was *stationary*. Whilst a subject is motionless within the environment their head movements represent their direction of gaze as they are looking around an environment. At these instances in time, a subject is not using their direction of gaze to navigate but to survey their environment.

It was decided to develop a method of graphical representation based upon this subset (only their head movements whilst paused) of the data. This new method, therefore, effectively combined both the pause points and direction of gaze data to produce a single unified representation. In this manner it is not only possible to determine where in an environment a person was pausing, but to discover in which direction they were looking whilst stationary. If, as it appears, subjects seem to be stopping at or in the vicinity of junctions, i.e. locations where a route choice decision needed to be made, then the ability to be able to reproduce and represent their head movements whilst making navigational decisions, might afford clues to this decision making process. The question can then be asked where are subjects looking in an environment in order to inform their decision-making?

An example of combined pause point and direction of gaze analysis is shown in figure 7.20 opposite. This figure opposite represents the path of a single individual navigating through World F.

The subject starts their journey in the top rightmost corner (on plan). As the subject enters the world, they initially remain motionless, turning to survey the new environment. They begin facing towards the world, turn around to look behind them (where they see the world boundary) and turn back through 180° again, before starting to move off.

The subject successfully negotiates two junctions before reaching a T-Junction. Here they are unsure whether to turn to their left or to their right. They pause to look in both directions. They elect to take the right path and continue their journey. Without hesitation they make a left-hand turn and rapidly reach a crossroads. Here they pause once more, looking to their left and right, before making the decision to move straight on. Continuing in the same direction, they walk past junctions on their left and right, before reaching another crossroads. Here they also pause to look around before deciding to move straight on as before. They finally reach the boundary of the world. This next part of their journey is not shown.

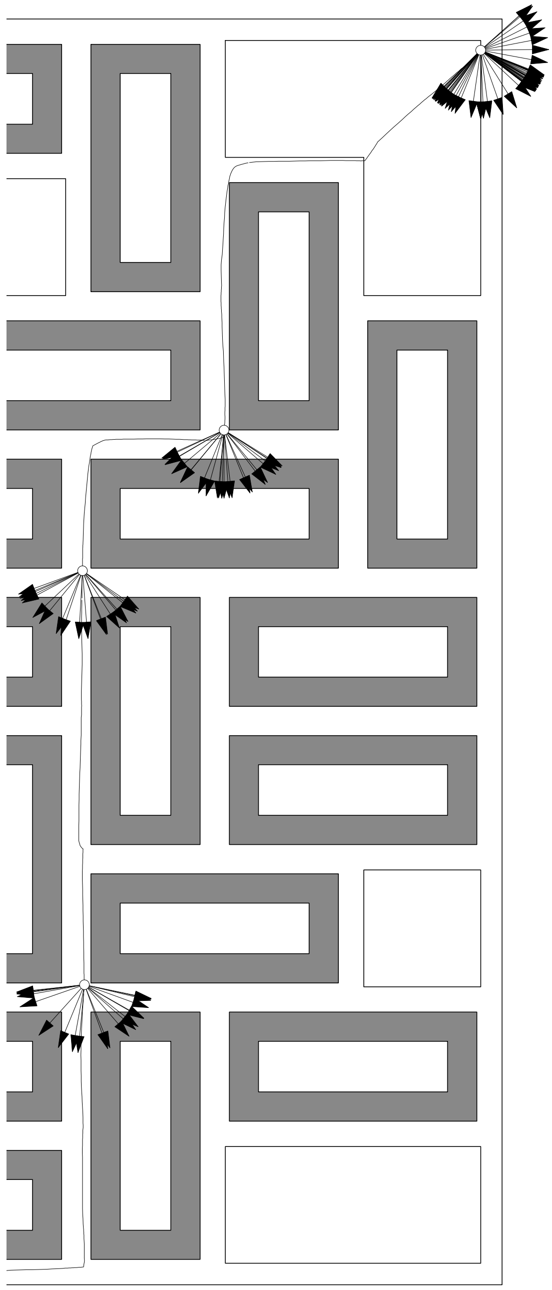

Figure 7.20 Path of an Individual in World F Illustrating their Pause Points and Direction of Gaze Arrows

The example of World C in figure 7.21 shows a person entering the world at the mid-left side. As in the previous example, they start by looking around. They do not pause again until they reach the main diagonal route to the monument. They pause a couple of times before reaching the monument. They pause just before the monument, where they stop, look at it and look around. On attempting to return to their starting position, they pause a couple more times, but fail to retrace their steps exactly.

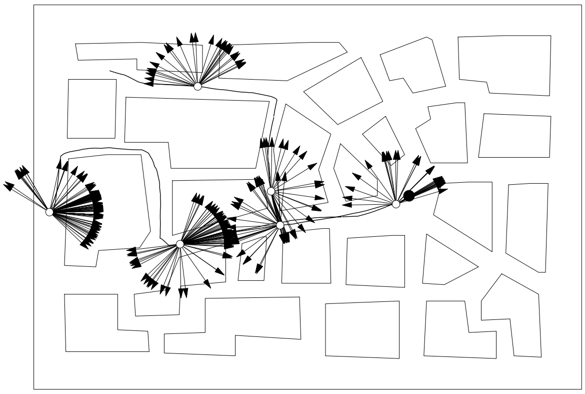

Figure 7.21 Path of Subject in World C

In World D opposite, the movement trace of another person can be seen. The positions where they pause and stop to look around are clearly discernible. This form of representation also contains an indication of time. The number of arrows radiating from a pause point is indicative of the amount of time that they are motionless. A pause point with only a couple of arrows radiating from it will represent a briefer pause than another location with a thick forest of arrows radiating from it. Another form of behaviour that is encapsulated in this form of representation can be clearly seen in figure 7.22 opposite.

Figure 7.22 Path of Subject in World D

Compare the upper-left pause point on plan (the end point of the journey) to the point on the far-left side (the start point of the journey). The radial distribution of the arrows represents the speed at which the person was turning around. If a subject spins around rapidly, the arrows are spaced far apart, such as the last pause point in the journey. If the same person is turning around very slowly that arrows will be quite densely packed. In summary, the number and spacing of the arrows radiating from a pause point can be interpreted as the relative duration of pause points and the speed that a subject turns around.

Conclusion

This chapter set out to examine a selection of smaller-scale actions that may take place during a journey. In particular, it considered locations in an environment where subjects pause for significant periods of time. Of all the methods of combined analysis and visualisation presented in this chapter, the final technique of illustrating pause points combined with approximate direction of gaze data appears to be a most successful way of representing the log file data and of capturing such small-scale actions. It is strongly suggested that where the subjects look whilst stationary is important. If this is true, then conversely what is available *to be seen* at any location must also be important. This leads to the suggestion that an attempt should be made to correlate movement data to both spatial and visual data. The type of data that most readily lends itself to such analysis is isovist data. Earlier in this chapter, it was conjectured that people appeared to be pausing in specific types of spatial location (junctions and centres of spaces). Therefore, the following chapter will begin to investigate whether there is a method of *characterising* where people are pausing in the environment in terms of that environment's visual, isovist properties. In the same way that isovist attributes can be calculated along a route generating a 'route vision profile' (Chapter 6), it is proposed to investigate the isovist characteristics of individual pause points.

Key Points

- Conclusion – People appear to be pausing at locations where route choice decisions need to be made.
- The k-means cluster algorithm can be effectively used to identify groups or clusters of pause points. However, the technique could be refined to include spatial information (such as mutual visibility between points), rendering it more useful for this task.
- A comparison of total journey time or distance travelled to stopping patterns may reveal individual characteristics of journeys.
- The combination of direction of gaze and pause point analysis appears to be the most useful method of combining the data and information available in a graphical form.
- There are possible differences between patterns of movement at the level of the building and at the urban environment level. These differences may be discerned through the duration of pause points.
- There is a possible correlation between the speed of movement and the duration of pause points. The faster a person moves, the less time they are likely to spend stationary.
- It is also possible that people pause in different types of locations in building 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.

Notes

1 This application was written in the C programming language.

2 Since a subject's position was measured ten times every second, in order to convert from the number of points sampled (see table 7.2) into the number of seconds spent stationary it is necessary to divide the number of points by a factor of 10.

3 Additional information exists on this sub-sample of five subjects, which might suggest reasons for the differences in their behaviour compared to the rest of the sample. By collating personal data provided by the subjects in questionnaires, it is possible to provide further information about the five subjects identified on the scattergram. These five subjects are all male. They have a slightly lower average age than the sample as a whole (28 years compared to 33 years) and an unusually high proportion of them were left-handed (three of these five subjects were left-handed).

4 Source is email correspondence with Dr. Charles Frederick Neveu, NASA Ames Research Center. Although John Stracham claims, again in email correspondence, that 7°[eye movement] is rare without head movement greater than 12°, citing the HMD handbook.

5 In conversations with Dr Graham at the Institute of Neurology, University of London, it was suggested that head movements captured by the headset used in these experiments were probably an accurate estimate of direction of gaze with an error factor of 10-20%.