A PATHFINDING APPROACH FOR REAL-WORLD SITUATIONS

Sarah Cook

MSc Artificial Intelligence

2010/2011

The candidate confirms that the work submitted is their own and the appropriate credit has been given where reference has been made to the work of others.

I understand that failure to attribute material which is obtained from another source may be considered as plagiarism.

(Signature of student)____________________________

Summary

Producing suitable routes for pedestrians in an urban environment poses many challenges, especially if the individuals concerned have cognitive impairments. This project offers a solution to some of these problems by constructing a system which employs principles taken from environmental psychology and spatial analysis, to produce a cognitive approach to the task. Initially, previous research is examined to determine the most relevant concepts and techniques in this area. From this, a number of system elements are developed, designed to approximate the main processes in human wayfinding. Simple models of artificial environments are constructed, and heuristics based on known route selection criteria are applied to these, producing elementary solutions. The results from this exercise show that the shortest path approach is outperformed by more cognitive alternatives, and that extending the psychological components of the system may create better solutions. A virtual user is then implemented, combining functionality to replicate user memory, the inclusion of landmarks, and the simulation of human wayfinding errors to both create and evaluate routes. These algorithms are used to investigate the effects of familiarity on other attributes of the solutions, and how a cognitive approach to evaluation can combine many of the problems encountered by wayfinders to assess the quality of the routes produced. Finally, the artificial environments are replaced by a model of a real site, and the behaviour of the system when applied to this is scrutinised. Despite the many complications which can be introduced by this type of representation the software functioned well, producing some high quality results.

Acknowledgements

Firstly I would like to thank my project supervisor, Roy Ruddle, for his support, advice and for helping to keep me calm throughout the whole process of completing this project. Also for his enthusiasm for the topic, and seemingly endless knowledge of the areas surrounding it.

Secondly my assessor, Marc de Kamps, for his feedback on both the interim report and at the progress meeting.

I would also like to thank my partner, Simon Hilton, for his endless help and support during this work and indeed my entire MSc. Without his constant supply of tea and ceaseless confidence in my abilities, I would never have completed it.

Finally I would like to thank my friends and family who have helped me celebrate my achievements and kept me going no matter how difficult it seemed to be.

Contents

Chapter 1

Introduction

With the increasing use of technology to provide navigational aids, and a rising percentage of the population with some form of cognitive impairment, the pressure to create suitable wayfinding tools for many different groups in society is growing. This project will focus on the requirements of pedestrian travellers, concentrating on the task involved in this type of path-finding and the techniques used to solve them. Unlike those using vehicular transport, where systems such as satellite navigation have traditionally been utilised, individuals travelling by foot have a far wider range of movement and very different wayfinding needs. Simply computing the shortest route to a destination may not produce the best solution, with factors such as complexity and understandability affecting speed and even the likelihood of successfully reaching the goal. These and other issues mean that alternative approaches must be considered.

The aim of this work is to produce a piece of software which uses a cognitive approach to find the most appropriate pedestrian route between two points, applicable to both individuals with loss of wayfinding skills, and those in new surroundings. As with human navigation, it will focus on the familiarity of an area, known or memorable landmarks, and other aspects that are considered during the formation of cognitive maps, rather than just adopting the Shortest Path approach. To achieve this, five main objectives were identified which correspond to the project's minimum requirements:

1. Determination of the cognitive principles to be applied to the problem.

2. Construction of a suitable model of the navigational environment.

3. Implementation of an appropriate route production algorithm.

4. Creation of suitable evaluation tools.

5. Establishment of an appropriate technique for data collection, and employment of this to gather relevant information.

This report will show that by examining previous research, a number of spatial analysis and environmental psychology principles relevant to the task of wayfinding can be chosen as the most influential on the resulting route. In addition, the cognitive issues surrounding how information is stored and deteriorates during navigation will be explored, and these concepts will be combined to form the basis of route construction and evaluation functionality. Once determined, the appropriate principles will be used to form basic heuristics, designed to create a series of routes, with each obeying one or more known human selection criteria. This method will allow different solutions with contrasting characteristics can be passed to an evaluation algorithm, and a virtual walkthrough technique will then used to determine the most appropriate.

The system itself was designed to be completed in multiple phases, and the report will indicate that by adopting this approach, the complexity of the resulting algorithms can be built up one element at a time. Beginning with several artificially generated environments, it will illustrate that the selected techniques and criteria are robust over a number of different models and with routes of varying lengths. With sufficient testing and evaluation at each stage, the different route attributes will emerge and be compared against appropriate baselines to show the performance of the application, and quality of the results. Lastly, the behaviour of the application when it is presented with a real environment will be examined. For the purposes of this project, the environment to be studied will be restricted to the University of Leeds campus, and a model which can represent these navigational conditions will be constructed such that it contains sufficient information about the domain without unnecessary complexity.

1.1 Project Methodology

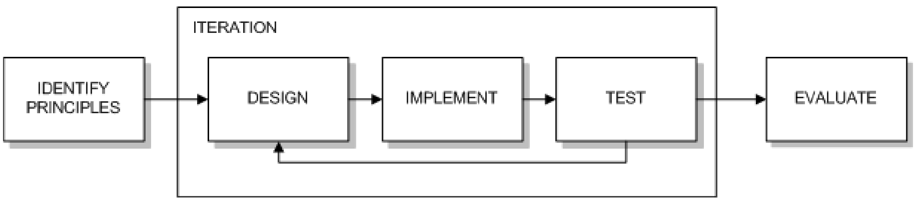

Figure 1.1: Project Flow

This project employed an iterative ‘design-implement-test’ methodology as shown in Figure 1.1 where, after the required principles had been identified, each cycle represents one system element corresponding to its associated phase of development. A plan of how the time spent on the project was divided between these iterations is given in appendix D.2. It illustrates that the work was broken down into three main tasks, with a long period of evaluation towards the end. Where multiple phases have been merged into a single time period, this indicates that aspects of one were required by the other, or that work was carried over with no significant break. This schedule varies in some parts compared to that proposed at the beginning of the project, and this will be discussed in section 6.3.

1.2 Report Layout

The remainder of this report will be divided into chapters in the following way:

Chapter 2 examines some of the many areas of previous research associated with wayfinding. It will explain how a suitable representation of the urban environment can be formed, determine the principles to be used for route creation and user simulation, and identify the appropriate techniques and metrics for evaluating the results.

In chapter 3 the system and its elements will be introduced, with details on how the complexity of the algorithms involved was increased as each phase of the development was implemented. The second section of this chapter explains how the basic environmental model was built, and how the initial heuristics were applied to this to create elementary routes.

The fourth chapter discusses the implementation of the virtual user, adding more cognitive principles and components to the system. It shows how user knowledge can be stored with suitable memory functionality, and examines the effect of extending the environment to include significant landmarks. Using these additions, it will indicate that more complex heuristics can be used to form adaptive routes through the domain. The remaining section in the chapter investigates how wayfinding errors and information loss can be combined with other psychological concepts to approximate certain aspects of the behaviour of a human navigating the suggested course. Using the results of this simulated walkthrough, a cognitive evaluation approach will be described.

Chapter 5 looks at the techniques employed to collect data from frequent users of the university campus, and the information resulting from this process. It will indicate that by scrutinising this data, the most commonly travelled areas of the campus can be identified for these individuals, and that this can be used to assist when converting the campus map into a suitable model of the site. Lastly, the performance of the system on this new representation will be examined, with the resulting routes being compared to those proposed by human wayfinders.

The final chapter discusses work attempted during the project which has not been reported elsewhere, draws conclusions about the performance of the system and quality of the results, and offers some guidance on the directions that future work could take. In addition, it indicates how the final project plan deviated from the original schedule, and why these changes were required.

Chapter 2

Background Research

2.1 The Problem

Unlike the approach taken by satellite navigation systems for vehicle route-finding, the intention of this project is to create a program which produces solutions that will assist individuals in the learning of routes, not just indicate the shortest path between two points. This requires combining basic wayfinding procedures with more cognitive techniques to produce the 'best' route for the user. In a society with an ageing population and a higher brain injury survival rate, the needs of people with cognitive impairments is becoming an increasingly important area of research. This project aims to produce an algorithm for providing assistance in the pedestrian navigation of both familiar and unfamiliar environments, which may form the basis of a solution for one of these needs.

Studies on both Alzheimer's and traumatic brain injury patients have shown that, despite significant cognitive deterioration in some cases, they can still solve many wayfinding problems. In Passini's work with Alzheimer's patients [30], the subjects' ability to navigate both familiar and semi-familiar environments was investigated. It was found that, with an appropriate number of well defined reference points, the majority of individuals could successfully locate destinations. Livingstone and Skelton [25] investigated a small group of high functioning traumatic brain injury survivors, and their performance when presented with a set of navigational tasks. These participants had difficulties in the absence of relevant landmarks, but their results improved significantly when these were introduced. Although two very different groups of test subjects, the wayfinding issues in both were found to be associated with the creation and use of cognitive maps. Solutions to overcome these obstacles have been proposed, but in most cases they require the installation of specialised equipment which may not be suitable for outdoor environments [5]. The method suggested by this project offers a far more adaptable approach, with no reliance on expensive apparatus or limitations with regard to setting.

Despite the current work being aimed at resolving the issues surrounding cognitive impairment, the resulting system will be equally useful to tourists or other visitors to a city or area. Many of the problems encountered in these situations are similar, including a lack of familiarity and a sense of disorientation when in a new environment. From this it can be seen that a solution for one challenge may also offer an answer to the other.

Producing a system capable of delivering this type of cognitive result is more complex than creating a simple route creation algorithm. It must include a variety of priciples taken from spatial analysis and environmental psychology, and combine these to produce an appropriate route between any two points. To achieve this, construction of the application must be broken down into separate tasks representing both the implementation and the evaluation of the system. Four main areas have been identified for this project:

1. Creating a representation of the environment to be investigated.
2. Producing an algorithm to generate suitable routes through this environment.
3. Construction of a user simulation for the selection of the best routes generated.
4. Identification of suitable techniques for evaluating suggested routes, including the collection of data about the environment which can provide a baseline for final testing of solutions.

These activities can be further broken down into subtasks, or phases in the development which relate directly to increasing levels of complexity, and each of them will be investigated in more detail in the following sections.

2.2 Representing the Environment

Research into the way that humans view their surroundings, and therefore construct routes, in urban environments has its roots in the pivotal work of Lynch in The Image of the City [26]. In this he examines the mental images formed by the inhabitants of three American cities, and how these affect their ability to navigate the world they live in. He identifies five specific elements which are contained within all cities, and used to define and structure our internal representation of the areas around us:

1. Paths – the channels through which pedestrians move. These could be pavements, roads, tracks or any other navigable feature.
2. Nodes – the decision points within our environment. A node may be an intersection of two or more paths, a point where there is a fork in the road or the location at which a user switches between transportation types.

3. Landmarks – visual references used by an individual to locate themselves within their surroundings. Landmarks are usually physical objects such as buildings or signposts, but may be local to or distant from the current position.

4. Districts – areas having a single, definable, characteristic. This characteristic may relate to visual, ethnic or class distinctions, or be a simple as the area which constitutes the main shopping area in a town or city.

5. Edges – barriers restricting or preventing pedestrian motion, or boundaries between different areas of the city. Unlike the first three element types, edges do not have to be physical but may be visual or cultural also. This category includes features such as walls, rivers and boundaries between districts.

By combining instances of these elements, any urban environment can be described in terms familiar to residents and visitors alike. Despite the importance of each of these features in forming a whole mental image of the environment, when applied to wayfinding they can be separated into three distinct layers of complexity.

The first of these layers can be considered the essential information needed to construct the most basic form of route. It consists of only the path and node elements, without which no route could be created. A path may consist of many segments or link, each joining two adjacent nodes. In a similar way, a route may contain many links or even entire paths, along with their connecting nodes, indicating how to travel from a start point to a destination point. Examining this layer more closely, we can see that it is sufficient to model basic wayfinding tasks such as the shortest path and travelling salesman problems.

The second layer of complexity includes the addition of landmarks, allowing the enrichment of the information provided to denote the paths and nodes within the route. Landmarks provide visual cues at navigation decisions, along paths or globally to reassure the observer of correct progress. Addition of this information provides a representation more closely resembling that produced when a person gives verbal directions for travelling between two points. Michon [29] discusses the role of landmarks in directions in some depth, and more details on this area will be given later. For most navigational tasks, this is the highest level of complexity that will be considered as it contains all the information required, and can be reduced to include few or no redundant or irrelevant details.

Adding the edge and district elements to the second layer provides the complete and most complex representation of the passage to be navigated. They give structure and heirarchy to both the environment and route descriptions. Districts signify a far coarser level of detail, encapsulating many examples of the four other elements. In some instances they may be used as the only description of how to travel to a destination, or even as the destination itself. An example of this are journeys between towns or countries, especially where the mode of transport is outside of the control of the traveller. Edges may also designate large features, but equally belong to the finer detail associated with paths and nodes. In the context of a University campus, a district may correspond to a department or even a playing field, and an edge the outside of a building or other such structure. These components will not be implemented within this project as they add little to the final model, but greatly increase the complexity.

Defining the essential elements within an environment is only the initial step in creating a representation for it. Finding a suitable way of simplifying these crucial components and their relationships is as important as identifying them in the first place. There are many alternative ways to approach this task, and only a few of them will be examined here.

2.2.1 Graphs, Grids and Deformed Grids

Graph Theory is a huge area of research centred around the use of nodes and edges[1] (links) to model relationships between entities. These connections may range from the spatial relations between two physical objects to syntax in language, the study of molecular structure, and beyond. For all these relationships, a common model format can be adopted as shown in Figure 2.1.

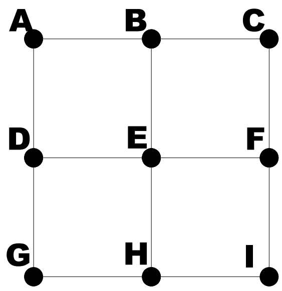

Figure 2.1: Example of Graph Format

Conversion of the paths and nodes in a basic environment to a graph or grid is one of the simplest and most successfully used forms of representation when considering wayfinding tasks. It is easy to see how swapping paths for edges to form graphs would be an attractive solution, requiring minimal expenditure of both time and effort. There is however a drawback to this approach. True graphs form perfect grids, with all edges having equal length and nodes being placed in symmetrical patterns. Unfortunately, very few urban environments correspond to this artificial layout, leading to the loss of all metric information relating to structure of the domain. In many route-finding tasks this is not an issue as merely the number of edges required to reach a destination gives a sufficient measure of success. Where this loss of data is unacceptable, graphs may be annotated with their cost (or length), but the use of a deformed grid [17] is a sensible alternative.

Deformed grids give a more realistic model of urban landscapes by employing irregular characteristics present in the environment when forming the representation, and are the first stage of conversion to a space syntax model. They show curves, lengths and the shape of buldings, restricting line-of-sight and pedestrian motion to the actual passages between structures. Metric information is retained in this

[1]These edges are very different to those described by Lynch [26], and signify only the existance of a relationship between the nodes.

way, whilst still producing a simple model. Deformed grids can be built by overlaying links and nodes on a map of the appropriate area, although as described by Werner [39] this process may have it's own difficulties. Figure 2.2 shows a very basic example of the type of representation which can be formed.

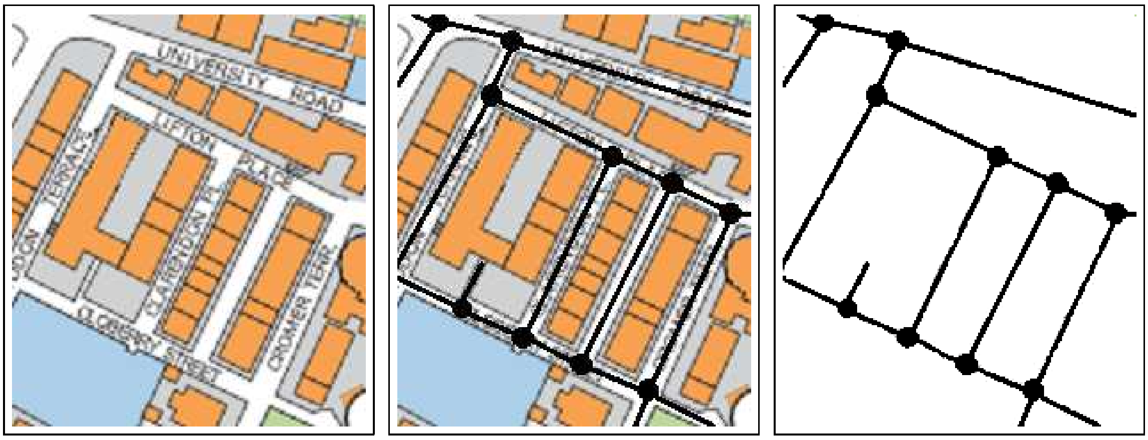

(a) Portion of Map (b) Partial Deformed Grid Overlaid on Map (c) Partial Deformed Grid

Figure 2.2: Example of How a Deformed Grid is Created

How the information contained in this representation is converted into a form suitable for computer applications is dependent on the problem being addressed, and will probably vary from author to author. In all cases a way of including the metric data depicted must be found.

2.2.2 Defining and Using Landmarks

Landmarks have an important role in the wayfinding success of individuals both with and without cognitive impairments. Experiments with patients with different causes and levels of cognitive deterioration [30, 25] have shown that they are crucial in the solution of navigational tasks, but tests on healthy subjects have also indicated similar relationships. Jansen-Osmann [20] used a virtual environment to investigate how children and adults use landmarks to navigate through a maze. This showed that visual cues significantly reduced the time taken to learn a route, and that more objects located at decision points are remembered than those placed elsewhere. Work conducted by Ruddle [32], also in a virtual environment, indicated that participants made fewer errors on routes with landmarks than identical examples with none. In addition, a difference between the success of using local or global landmarks was also observed by this study.

A local landmark is considered to be one which is only visible in a restricted area and may indicate either the position or position and direction of a required turn in the route. They provide positional information to the user, and are relied on heavily during navigation. In contrast, global landmarks are distant from the route, but usually visible from many locations. They allow the observer to orient themselves in the environment, but are likely to be less important for wayfinding tasks. Ruddle [32] suggests that only local landmarks are required to successfully navigate, although it is pointed out that this may be due to the environment and approach selected.

From Hurlebaus [18], an alternative hypothesis can be formed; that different individuals use different landmarks. Lynch [26] suggests that the selection of relevant visual cues may also be affected by a person's familiarity with a location, with smaller objects being chosen by those with increased exposure. Furthermore, according to Allen [2] and Sorrows [34], different wayfinding tasks may require specific types of landmarks such as the use of objects with visual or cognitive rather than structural characteristics to navigate familiar routes. These variations illustrate the difficulties which may be encountered when trying to determine a 'good' landmark in a given situation. The following list is a selection of attributes, compiled from those given by a number of different authors [34, 35, 12], which should be considered when identifying suitable objects. By combining as many of these characteristics as possible, the effects of the problems described can be minimised.

- Singularity – A visual contrast with its surroundings making it conspicuous.
- Prominence – A location which is at an important junction or visible from many locations.
- Accessibility – A location at which many paths converge.
- Content – Cultural significance or a distinguishable use or meaning.
- Prototypicality – Good representation of a category, or how different it is from other members of the same class.
- Reliability / Persistence – A high likelihood that this landmark will exist at a later date and that a user will be able to find it.
- Relevance – A high importance for navigation.

Once defined, only the task of integrating landmarks into the environmental model remains. There is no standard approach to this, with methods varying from the designation of whole areas of influence, to simplification to a single point. Given the types and size of landmarks within the campus site, it should be sufficient to use the latter more straightforward technique, although the point size will be allowed to vary with the magnitude of the object or structure.

2.2.3 Alternative Representations

Although one of the most obvious and simplest, using Lynch's [26] elements to create a deformed grid is not the only possible representation for an urban environment. An alternative is the approach taken by space syntax [17], a relatively recent field of research. In this, axial maps are constructed from a deformed grid by drawing the fewest, longest lines through the grid in order to cover it, and axial grids by representing nodes with lines and edges with line intersections. By combining these two models, measures of the connectivity and integration of routes can be found, and used to determine the likelihood that it will be selected by an individual. Despite the merits of this approach, it can be complex and allows only one wayfinding heuristic to be tested.

Another example of the type of representation possible, is the use of first order or predicate logic. This method generalises the content of the environment to objects and their relationships, along with rules pertaining to how these can be combined or the actions associated with them. The logical approach has led to many ontologies such as the one suggested by Kikiras [22] and even entire models [15, 23]. It is best suited to simulating process such as storage of route information or selection of predefined routes, but can also be applied to other decision making systems.

These and others produce structurally very different outcomes, but they may all be considered extensions to the original representation or variations on the same theme. All start with some or all of the elements described by Lynch [26], and differ only in how these are converted into a usable form.

2.3 Routes

Once a representation of the environment has been formed, the task of using this map to create and travel along a route is the next logical step. This is termed wayfinding, and is described by Allen as:

> *"Purposeful movement to a specific destination that is distal and, thus, cannot be perceived directly by the traveller."* [2, p.47]

The success or failure of a wayfinding solution is determined by the ability of the individual to navigate between start and destination points, within a limited time and appropriate spatial distance. In general, the task can be subdivided into three possible categories stated by Allen [1] as:

- 'Commute' wayfinding – the travel between two known points via a familiar route.
- 'Quest' wayfinding – travel from a familiar location to a previously unvisited point, where the information on the destination is provided by maps or verbal directions.
- 'Explore' wayfinding – investigation of the surrounding environment by travelling into an unknown area.

No consideration will be given here to the exploratory wayfinding task as it is beyond the scope of the project, but the remaining categories will be present in one or more of the elements to be produced.

2.3.1 Selecting a Route Between Two Known Locations

By its very nature, an urban environment is dynamic and forever changing habitat, with new structures being constructed and old ones dismantled almost continuously. In addition, events may close roads or even force whole sections of towns and cities to be cordoned off. Notwithstanding these ceaseless changes, the sheer number of possible routes through an environment may make the problem of selecting a single one seen almost futile. Despite this, most people and many non-human species complete this task with very little conscious thought. There are several schools of thought about how this is achieved, but the most widely held suggests that humans use a basic set of rules and attributes, known as route criteria, to create a single solution, adding small detours only where absolutely necssary.

2.3.1.1 Route Criteria

There are many criteria thought to be used by humans when choosing the most appropriate route between two known locations, with variety of factors which affecting the decision, from those of distance and complexity to more subjective characteristics [13]. Golledge [14] established that of these, the ten shown in Table 2.1 can be defined as the most commonly used.

Rank	**Criteria**	**Description**
1	Shortest path	Minimises the physical distance between the start and end points
2	Least time	Minimises the time taken to travel between the start and end points
3	Fewest turns	Minimises the number of changes in direction between start and end points
4	Most aesthetic	Routes near to parks or other public spaces, and away from waste disposal or similar sites
5	First noticed	Minimises the time taken to identify
6	Longest leg first	Longest segment without turns is the first
7	Many curves	More curved or partially curved segments
8	Many turns	Maximises the number of changes in direction between start and end points
9	Different from previous	Not previously or recently travelled
10	Shortest leg first	Shortest segment without turns is the first

Table 2.1: Route selection criteria.

Each of these criteria correspond to a single attribute of the route, but whether through coincidence or design, one route may fit several criteria. For example, if it can be assumed that a constant speed is maintained throughout all journeys, it is clear that the first two criteria will result in identical routes. Conversely, attributes such as aesthetics, previous travel and the time taken to identify are likely to specific to the individual, and may produce many alternative solutions for the same navigational task. In addition, Golledge also found that different approaches may be chosen depending on the direction of travel, with variations between inbound and outbound routes [14]. This may be due to the real or percieved configuration of the environment, or just personal choice. The purpose of the trip, and whether the trip has a single or multiple destinations, can also affect the route chosen [13]. For everyday journeys such as travelling to or from work, the focus is likely to be on the time taken to reach the goal, whereas leisure travel is more likely to be influenced by the scenery surrounding the route. Attempting to include all of these factors when selecting a suitable route would create a highly complex system, often producing results which may be contradictory or confusing. To avoid this, only a small number of these criteria should be chosen to be used separately or combined. For this project, shortest path, fewest turns and longest leg first were chosen as these are some of the most objective solutions, and through the use of a well known algorithm, heuristics which can construct routes conforming to these criteria were then formed.

2.3.1.2 Sticking to Familiar Areas

There is one significance absence from Golledge's list, that of familiarity, which is likely to be due to the method of testing[2]. Allen [1] indicates that habitual locomotion is the most frequently used navigational technique for travelling between two know locations, and that its employment can become almost automatic. How experiences are used to form mental maps will be explained in section 2.4.1, but here it will be sufficient to say that this process can produce well remembered, if somewhat disconnected, areas in only a short space of time. As this familiarity is built up with repeated exposure it begins to influence a wayfinder's decisions, with travel being restricted to a small number of 'major' and well known routes [24]. This behaviour has been linked to the configuration of features within the environment [24, 17], the use of boundary relations [24] and simply the effects of habitual performance reinforced by memory [13], but whatever the cause it is a frequently observed phenomenon. When creating a cognitive solution, careful consideration should be given to the impact that user familiarity has on route choice, integrating some measure of this attribute into the techniques employed.

2.3.2 Creating a Route To a Novel Destination or Through an Unknown Environment

Travel through an unknown or partially known environment poses problems in addition to those already described. Exploration is a technique that is time consuming and has no guarantee of success in situations where start and goal points are different. To avoid this, a suitable source of information about the area to be traversed must be found. With the advent of the internet and increasing use of visualisation software, this knowledge may come from a variety of sources and may be displayed in a number of different forms, but they mostly fall into the two broad categories of maps and directions.

2.3.2.1 Maps and Directions

A map is any two or three dimensional graphical representation of a route or environment, from a simple line drawn on a page, to a complex road plan or atlas. The use of maps is thought to have originated up to 8000 years ago, with evidence of the depiction of settlements coming from early cave drawings. In more recent times, these representations have progressed from static two dimen-

[2] Golledge used maps of realitively unknown areas in his experiments, but these can be considered as representing known locations as the entire environment is depicted in the map. The overall effect is to simulate how the route was first selected, but once chosen its structure is unlikely to alter significantly over time.

sional illustrations to three dimensional and interactive models. Whatever the format, the purpose of transferring environmental knowledge from one person to another remains.

Conversely, directions use spoken or written words to indicate how to travel from one point to another. They are usually sequences of locations and actions describing the traversal of a route, rather than a general depiction of the environment. Studies have shown that although their content may vary from person to person, to be useful directions must encompass a core set of essential details [9, 8].

Research by Meilinger [27], in which participants learnt two routes through an urban environment from either a map or verbal directions, shows that the resulting wayfinding performance is equal for both of the approaches. Indeed, Tversky [38] not only suggests that route depictions and directions are interchangeable, but also proposes a that there are a set of protocols that enable automatic translation between the two. This is surprising as the different tools contain disparate types of information. A map represents survey knowledge incorporating metric and structural relationships, whereas directions consist route knowledge made up of the sequence of actions and locations required to be performed in order to reach the goal. Despite there appearing to be little separating the performance of both approaches, each has its own strengths and weaknesses and these should be considered when choosing the best method for any task.

Environment maps contain a large amount of information about not only the route, but also the features surrounding it. They allow the user to select landmarks which are already known to assist with navigation, or impose their own personal preferences or informational needs on the solution. Unfortunately, this abundance of data can lead to this form or representation being complex and difficult to read [37], with individuals being easily distracted by large features. In contrast, route maps contain only the details required to successfully navigate between two points, or those pertaining to the region immediately surrounding the passage to be used. They significantly reduce the complexity of the representation, but also lose many of the advantages associated with this form of description.

Directions are similar to route maps in that they provide only information about the nodes and paths to be traversed, combined with visual cues along the way. They contain little or no reference to the extended environment or areas not visible from the route. This makes them very specific to the individual journey, but this has its own advantage. By limiting the description, only those details essential for navigation are transferred to the wayfinder, reducing the complexity and the amount of information to be remembered. In contrast, the lack of further knowledge may mean that travellers cannot correct errors or use distant landmarks to reorientate themselves if lost.

One of the most important factors affecting the usability of any route description is the reliability of its source. Problems with externalising an internal representation can be encountered in either approach, and these will be examined when methods of data collection are investigated later in this report (see section 2.5.2).

2.3.2.2 Wayfinding Strategies

With sufficient information gathered, a user is ready to proceed with the task of navigating to the destination. In order to aid in this process, one of a number of wayfinding strategies [2] can be selected. The following list considers only those approaches relevant to either ‘quest or ‘explore’ tasks, and indicates which strategies are appropriate for each.

- Oriented search – A user orients themselves in the expected direction of the goal, and then searches for the destination. This strategy is suitable for both types of wayfinding.
- Following a marked trail – This could be following a series of signs, colour-coded trails, or even a rope or tether. As with the previous strategy, this would also be appropriate for both types of travel.
- Piloting between landmarks – A user relies on a sequence of landmarks and actions which, if followed in order, give a method for locating the destination. This strategy is also suitable for both ‘quest’ and ‘explore’ wayfinding.
- Path integration – This strategy relies on the continuous monitoring of location and orientation through the use of external cues such as landmarks. Although this could be used for any wayfinding task, it is most suited to exploratory travel.
- Referring to a cognitive map – Use of a pre-existing internal representation of the environment to find new or known locations. It can be used for both types of wayfinding, but requires sufficient knowledge about the setting.

Of these the latter three are most commonly employed, but an individual is far more likely to combine aspects of each of them when navigating through an environment, rather than rely solely on one. All of the described strategies depend on the traveller possessing certain cognitive skills, and being able to call upon them as the journey progresses. Amongst these are the ability to recognise locations through landmarks or structural features, to monitor movement and speed, and to remember previous locations. In addition they may each have their own complexities such as identifying suitable visual cues, or be error-prone due to a lack of experience.

For this project, the most suitable strategies centre around the use of landmarks and internal representations. These permit controlled simulation of a human wayfinder, and will be discussed in more detail in the following section.

2.4 Simulating Human Wayfinding

When simulating human wayfinding, more the just the locomotion between two points must be considered. This section will examine how information is stored during the navigation process, how a simple strategy can be employed to traverse an environment, and how distractions and momentary lapses of concentration can have an adverse affect on performance. A number of different wayfinding strategies were described in section 2.3.2.2, but one of the most widely used, which has been selected for implementation in this simulation, will be examined in more detail here, with its components and concepts explained.

2.4.1 Cognitive Maps

A cognitive map is an internal representation of information about an external environment, formed by gathering knowledge through either examining maps or actually navigating parts of it. The process of constructing this is gradual, with the mental model being continuously updated and extended as an individual experiences more and more of their surroundings. This cognitive representation consists of geometric information such as points. lines and areas, but may also contain details about meaning, use and other non-metric relationships. It is thought [33] that there are three types of information used to encode a cognitive map, with data passing from one to the next as a person's familiarity with an area increases.

1. *Landmark knowledge.* This is the initial type of information stored in a cognitive map, and is usually formed quickly through only brief exposure to the environment. It relies on an individual simply identifying and remembering prominent and easily recognisable landmarks in the locality, with no regard as to their actual position or connection to other features. These landmarks may be visual cues at the start and end of the journey, with a few of the most conspicuous structures interpersed between the two. By comitting these to memory a traveller will have an indication that the destination has been reached, or that they are at least travelling in the right direction. There is some debate over whether this is truly the first stage in forming a mental representation, but early tests on children [33] indicated that it may be considered separate for this section of the population at least.

2. *Route knowledge.* Once a number of significant landmarks have been identified, increased navigational experience allows these visual cues to be linked into sequences showing the order in which they are encountered throughout a journey. By combining this with the actions to be performed at their locations, a procedure to navigate between two points is formed. At their bare minimum, these procedures must identify the points and actions pertaining to changes direction required to progress on the route, but may also contain impressions of distance, angle and terrain. This type of knowledge provides shape to the representation, and as the number of known routes increases, the stored information progresses to the final type of knowledge.

3. *Survey knowledge.* By identifying points of intersection between routes, the individual can connect these separate pieces of knowledge into a whole, enabling them to adopt a single frame of reference. Through the addition of directional headings to this, survey knowledge is created. This encodes topographical properties into the representation, allowing objects to be positioned with respect to a fixed coordinate sytem, establishing the shape of features and the euclidean

distances between them. Determining these relationships enables information about a location to be retrieved without first having to mentally traverse a route to reach it. Although the addition of metric information may seem to suggest that this knowledge would resemble a physical map, there are some indications that this is not the case, with it containing data from a number of senses, not only the eyes.

Experiments [36] have shown that the use of physical maps may bypass part of this process, allowing information to be stored directly as survey knowledge. However, later work [38] indicated that, when faced with a completely unknown environment, individuals form sequences of locations and actions even when supplied with data from this type of source.

Cognitive maps can be used either as a wayfinding tool, or to organise the spatial relations between objects or locations [13]. Although they can encode both metric and non-metric information, they are rarely accurate or complete. Lack of detail, missing areas and mental distortion can make them error prone, but repeated exposure to the environment can lessen the effects of these problems. Despite their shortcomings, humans tend to rely on cognitive maps far more than any other type of wayfinding strategy, indicating their importance. For the application proposed here, this concept of internal representation will be used mainly in the storage of information about the environment, rather than assisting with navigation through it. Despite this, its content and how this information can be used to create a cognitive route will be considered.

2.4.2 Piloting Between Landmarks

Piloting between landmarks is one of the simplest wayfinding strategies that an individual can employ to navigate between two points. It requires no prior knowledge of the environment or complicated monitoring of speed and orientation to be successful, although some information about the route to be followed must be supplied. These attributes, combined with its applicability to any type of wayfinding task [1], make it a very common approach. In order to accurately pilot between landmarks, the information gathered about the route must be converted into a sequence of location-action pairs, where the action is the movement required to travel from the current location to the next in the sequence. By completing this series of steps, a traveller should be able to navigate from the required start point to the specified destination. In human wayfinding, the number of landmarks identified and the relationships between them increases with familiarity, resulting in a full sequence of those required to navigate the complete route. From this description it is clear that this information is comparable with either verbal or written directions designed to transfer the essential details of a journey from person to person, or route knowledge stored in the cognitive map (see previous section). This relationship indicates why this particular wayfinding strategy is appropriate in such a wide range of situations, and why it has been selected for implementation in this project where both sources of information are present. To be usable, piloting requires the possession of three main cognitive abilities; object recognition, paired associate-learning and sequential learning. These concepts and abilities will be modelled directly through the use of a user simulator, with increased recognition and familiarity leading to fewer errors, conversion of directions into location-action sequences, and these sequences being used to 'walk' a virtual route.

2.4.3 Navigational Errors

Even when travelling between very well known locations using frequently traversed routes, people regularly make minor everyday wayfinding errors due to distractions or momentary lapses in concentration. These mistakes are usually easily corrected, and only affect a traveller's behaviour and ability to reach a destination in relatively insignificant ways. However, if the individual concerned is easily disoriented or confused, as is the case with many cognitively impaired people, these deviations may pose a very serious problem. In order to mimic the user successfully, these lapses must be examined in detail and the causes or effects incorporated into the user simulation algorithm.

From studying a group of 29 participants for four weeks, Williamson [40] defined nine types of errors which were regularly encountered by the test subjects:

- Wrong turning – turned one way when should have turned another.
- Missed turning – travelling past turning.
- Route selection error – inability to make a decision about the correct route.
- Misconception of location – not where expected to be.
- Travelled to incorrect location – ending up somewhere other than where intended.
- Premature exit – believing further on along route than actually are.
- Return route error – forgetting how to get back.
- Route exit failure – continuing on a journey section when it should have ended.
- Miscellaneous – any errors not covered by the previous eight categories.

Of these, making a wrong turn or failing to turn at all accounted for 49% of the maistakes made, indicating that these were by far the most common errors. Each of the categories were experienced in both familiar and unfamiliar environments, and whether travelling by foot or utilising another form of transport. This shows that all types of wayfinding task are likely to encounter some loss of performance due to these types of mistakes.

Many different reasons were given for the navigational lapses, but they all fall into one of five categories; environmental cause, inattention, inadequate knowledge, habit, or inadequate cognitve map. By further breaking down the first class into physical and psychological reasons, it was shown that 82% of errors resulted from cognitive causes. As this project seeks to evaluate routes from a cognitive standpoint, the inclusion of wayfinding errors as part of this process would benefit the outcome greatly. Return route error is beyond the scope of the user simulator as it is designed to travel in only one direction, and there will no consideration of the miscellaneous error category as it is not apparent how this could be replicated. With these two exceptions, all other errors will be incorporated, along with the associated frequencies, [40], into the simulation algorithm during the early stages of development.

2.4.3.1 The concept of 'forgetting'

In addition to the errors that occur during navigation, problems with wayfinding can arise before a single step has been taken. Deterioration in the data stored in a cognitive map and loss of information through partially forgotten directions, can have a significant impact on a person's ability to reach their destination, or even know which way to go. The dictionary definition of the word forget is as follows:

> *"Forget: To fail to remember (someone or something once known), to neglect, either by mistake or on purpose.*" [6]

An alternative and slightly more appropriate description is that forgetting is the deterioration of knowledge, usually due to the passage of time. There have been very few studies on the actual effects of forgetting when applied to wayfinding, but we all know that it exists and are likely to have experienced it many times in our own lives. What research does exist relates mostly to the acquisition and retention of landmark information, and how this affects an individual's ability to successfully navigate a route. Janzen [21] establishes that objects occuring at decision points are more likely to be remembered than those elsewhere, and Stankiewicz [35] suggests that the details of structural and distinctive landmarks are acquired faster and retained for longer. A further study, conducted by Corazzini [7] investigated the retention of route versus survey knowledge by testing individuals' ability to navigate a virtual environment after different periods of time. The most important finding of this work is that route knowledge is more prone to loss than survey knowledge, probably due to its rigid structure. It also observes that the names of streets tend to be the first information to be forgotten. In both papers, the familiarity of the environments involved is given as one of the main factors affecting both acquisition and retention of route features. As individuals' exposure to both routes and environments increases, the amount of detail remembered rises and the type of information retained changes.

Using this information, a general model of forgetting can be formed. By considering the familiarity of a route component either stored in memory or given in a set of directions, a judgement can be made as to whether or not it should be forgotten with more familiar examples retained. When selected to decay, different elements of the information will be removed at different rates, according to their function or role in the representation. Using a simple set of rules, concepts such as the loss of specific details before route structure is degraded can be managed and used to partially replicate the effects seen in humans.

2.4.4 Previous Models

There have been a number of attempts to model the concepts surrounding human wayfinding, although the majority focus on the acquisition and retention of spatial knowledge rather than the process as a whole. An early example of this is the TOUR model developed by Kuipers [23] to simulate the storage of information in a cognitive map. This system uses predicate logic to define objects and their relationships, along with a series of rules to control movement, detection of features and other such objects. By replicating the motion of a user through a 'You Are Here' pointer, knowledge is assimilated, stored and translated for use in later tasks. Although somewhat successful, this model is complex and has many limitations such as its lack of errors or loss of data.

In a later model, Gopal [15] employs two separate modules to mimic the process of initially creating, and then storing, a representation of a surburban environment. The first of these modules uses both global and environmental frames of reference to create an objective model of the surroundings, taking inputs such as the orientation of streets and characteristics of other objects and features. It simulates the movement of an individual through these, translating each piece of information into a predicate logic form. The second module then recreates a realistic process of filtering this data and storing it initially in working memory, but later transferring it to long-term memory. For this, only the most salient objects are selected to be 'remembered', with the rest being removed by object and scene filters. Once in the working memory, the information is further deteriorated by interference and decay, having the net effect that very little is transferred to the long-term memory. By adjusting the values of several external variables this model can reproduce the differences in ability found between individuals, both for creating and storing representations. This approach offers a more realistic model of the cognitive process than the previous TOUR example, addressing the main problems described.

Both of these models show that cognitive elements can be successfully simulated through the use of computer algorithms, but they only replicate part of the wayfinding process. No model which incorporates so many different aspects as the application proposed here could be found, eliminating the possibility of direct comparison.

2.5 Evaluation Techniques

Before deciding how to gauge the outcome of an application, it is important to determine what characteristics of the system need to be tested. Is it sufficient to examine the speed of computation and other properties of the algorithms, or are the solutions produced of greater significance? The simplest way to evaluate any system is to select and compare one or more performance metrics, and this section looks at how these can be chosen, and how relevant data can be collected where appropriate.

2.5.1 Wayfinding Metrics

The purpose of any metric is to provide a single value indicating some measure of the performance of a system, and although the time taken to produce a solution is of some relevance when judging the success of a route creation algorithm, the emphasis is definitely on the quality of its output. With this in mind, only techniques for examining the resulting routes will be investigated in this report. In wayfinding there are several separate areas which can be tested, and a number of different approaches which can be taken to gather the required results. For the proposed application, the most suitable evaluation criteria fall into two broad categories, those associated with the route itself and those linked to the behaviour of a user navigating it.

When evaluating the suitability of a route, many aspects of its structure and useability may be relevant route metrics. Attributes such as length, complexity and intelligibility can all affect the success of an individual traversing it. A good place to start when trying to identify ways of measuring this ‘performance’, is the route criteria given in section 2.3.1.1. By basing metrics on these characteristics it is possible to directly measure how attractive a route is to a would-be user, giving an indication of its suitability. As discussed previously, some of the attributes involved in these criteria are subjective making them difficult to quantify, but others such as length and number of turns are physical, requiring only simple computation to calculate them. Through the selection of appropriate baselines from either pre-computed thresholds or specifically generated routes, these values can easily be converted to metrics showing how good a particular route is.

A further measure of route quality can be found by examining the effects of familiarity when both selecting a route (section 2.3.1.2) and navigating it (section 2.4.1). Increasing this characteristic may significantly improve success, so a solution containing many well known components would be more advantageous than one with none. Finding a suitable way of normalising the resulting values for this may be difficult, but the benefits of appraising the level of this attribute may outweigh the issues it poses.

An alternative approach to evaluation for this type of metric is to compare the resulting solution with a route which has been supplied by a human user for the same start and end points. If this user data is taken to be the ‘perfect’ method for navigating between the two required locations, then the resulting measures will be similiar to the ’perfect search metric’ described by Ruddle [31]. In addition to scrutinising the differences between the supplied and generated routes, examining the degree of correlation in the selected components may give another measure of suitability.

This second category of wayfinding metric relates to the performance of a user navigating a route rather than the route itself. According to Ruddle [31], it can be subdivided further giving three subcategories covering task performance, physical behaviour and cognitive rationale. The latter two of these involve monitoring body movement, classifying errors or investigating cognitive reasoning and will not be detailed here, but the third, that of task performance, is applicable to the success of either a human or virtual user, warranting its inclusion. Time taken, distance travelled and number of errors made whilst completing a navigational task all fall into this subclass of metric along with other measures concerning distance to the first error and similiar observations. Comparison of these values between test subjects can illustrate individual differences, and between routes highlight difficulties traversing them. Careful consideration must be given to the effects of external factors when employing this type of metric however, ensuring that issues such as an individual's normal travelling speed do not distort the results. A major drawback to this type of evaluation is the need to find a sufficient number of test subjects, although if suitable simulation functionality is available, it may be possible to use emulated behaviour rather than actual humans.

2.5.2 Collecting User Data

Whether gathering suggested routes or attempting to collate task performance results, retrieving suitable information from users poses a significant challenge. Data collection in this field may encounter several problems, the most serious being related to the transformation of an internal representation into an external recordable form. Although a person may know exactly how they travel between two points, or the difficulties they experienced, in order to pass this information on to another person they must elucidate it first. This section will examine some of the techniques available, in each case explaining their strengths and weaknesses.

One of the most obvious ways of collecting wayfinding data is through the drawing of maps or plotting of routes, whereby an individual produces a graphical indication of the environment and required movement. Of the two approaches, sketch maps are probably the most useful, normally showing the most prominent features and those elements essential to navigation [26]. By drawing the paths, landmarks and other components of a journey, the spatial relations governing them can be identified. This provides a representation of not only the features required to reach the destination, but also an impression of the distance between them. Unfortunately, this method is very prone to error, arising from distortion and fragmentation within a person's cognitive depiction [26] in addition to their ability to adequately sketch this knowledge. In a similar way, the seemingly easier approach of marking a route on an existing map requires that the user can fully understand the graphical representation and locate themselves within it [37].

An alternative method for communicating environmental information is through verbal or written descriptions or directions. Route descriptions are usually route maps in verbalised forms (see section 2.3.2.1), and are regularly presented alongside this type of representation. They frequently give details of the environment surrounding the course of travel, and in many cases assume knowledge of the use or significance of certain smaller features. On the other hand, directions normally involve providing information for people who have little or no familiarity of the environment in question. For this, much larger and more distinctive landmarks are likely to be chosen along with main thoroughfares rather than complicated shortcuts. A drawback to this approach is that in order to give a set of directions or description of a route, it requires that the individual providing them have experience of journeying between the start and destination points. According to Denis [9], a useful set of directions must contain a number of core elements identifying locations and actions crucial to successful navigation. It is possible that the person describing the route has not travelled it in some time, meaning that they must call on their mental map and perform the corresponding calculations. This process can produce errors, resulting in missing or inaccurate instructions [8, 9]. In addition, many individuals elaborate on irrelevant details, increasing the complexity of the depiction unneccessarily and creating sources of confusion. Despite this, using a verbal or written portrayal reduces the distortions found in sketch maps by removing the need to give distance and angle information, and eliminates the need for graphical abilities (a typing error is far easier to correct than a misplaced path on a map).

Having established the type of depiction to be collected, only the question of how to gather this information remains. Employing an external observer, using interviews or questionnaires, and collating the data in user diaries are the three main options. The first of these requires that a tester 'stalk' the individual, and although providing very accurate data this has been shown to affect the behaviour it is supposed to observe (see Zimmerman [42] and Hill [16] for discussions on this). Both of the two remaining methods are known to approach the accuracy of the first, and have been successfully employed in several studies where they were selected for different reasons depending on the aims of the authors. Two examples of differing techniques are Williamson [40] where diaries in which subjects record navigational errors are used, and Denis [9] and Michon [29] where sets of directions are collected through interviews to examine their usefulness for wayfinding. Although the approaches are very similar, they provide slightly varying details such as the type and size of landmarks.

In wayfinding, producing a diary is a simple technique whereby individuals are asked to record their movements for a short period of time, giving details of how they navigate everyday routes. They result in descriptions of areas both familiar and unfamiliar to the subject, frequently using less obvious landmarks and background knowledge of the areas being traversed. In contrast, interviews may produce more generalised descriptions or just basic sets of directions, depending on the type of question asked. Most interview and questionnaire approaches require that the subject has knowledge of the area being discussed, although they may be carried out after an experiment designed to form this. In general, which collection technique is selected will likely depend on the resources available, the specific task being investigated, and the required level of accuracy.

Chapter 3

Introducing the System and the Environment

3.1 Overview

As the aim of this project is the production of a cognitive solution to the wayfinding problem, it was clear that this application would consist of more than a simple route creation routine. With this in mind, the system in Figure 3.1 was designed and constructed.

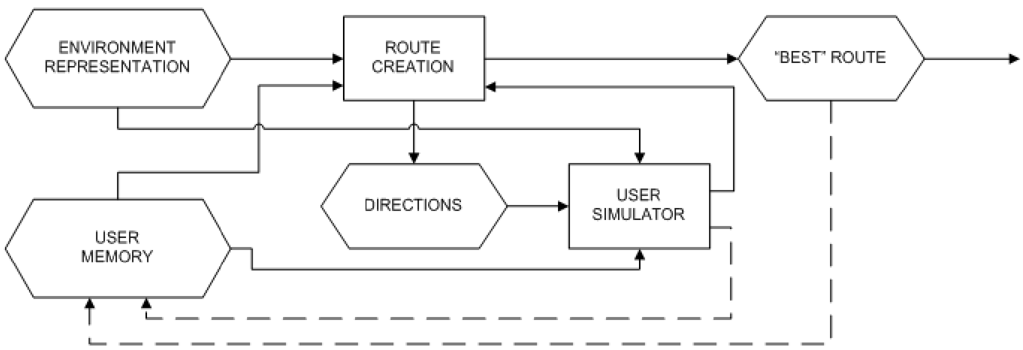

Figure 3.1: Elements of the system

This algorithm combines data input and storage facilities with route creation and evaluation functionality, incorporating cognitive principles throughout not only the construction process, but also that of selection. By breaking the system down into separate elements, it could be built and tested in multiple phases, with additions and modifications made at each stage. Table 3.1 shows how the complexity of the application is increased with each step, incorporating more elements or changing existing ones.

Using this approach allowed each component to go through many subphases of development, and the values for the variables investigated where possible. All programming has been completed in C++ as it is the most familiar language to the author, and allows for the algorithm to be constructed in separate stages corresponding to the phases of development. In addition, only the most basic standard libraries were used in order to allow simple transition between platforms (the initial stages were created on a machine with a Windows operating system, but the final product must be run on a Linux based system).

<table>
<tr><th>Phase</th><th>Basic Environment (Phase 1)</th><th>User Memory (Phase 2)</th><th>Landmarks (Phase 3)</th><th>User Simulation (Phase 4)</th><th>Real World (Phase 5)</th></tr>
<tr><td>Environment</td><td colspan="2">Basic Grid</td><td colspan="2">Sufficient Grid</td><td>Real</td></tr>
<tr><td rowspan="4">System Elements</td><td colspan="5">Environment Representation</td></tr>
<tr><td colspan="5">Route Creation</td></tr>
<tr><td></td><td colspan="4">User Memory</td></tr>
<tr><td></td><td></td><td></td><td colspan="2">User Simulation</td></tr>
<tr><td rowspan="6">Route Creation Heuristics</td><td>Shortest Path</td><td colspan="2" rowspan="3"></td><td colspan="2">Shortest Path</td></tr>
<tr><td>Minimum Turns</td><td colspan="2">Minimum Turns</td></tr>
<tr><td>Longest Leg First</td><td colspan="2">Longest Leg First</td></tr>
<tr><td rowspan="3"></td><td colspan="4">Shortest Path + Familiarity</td></tr>
<tr><td colspan="4">Minimum Turns + Familiarity</td></tr>
<tr><td colspan="4">Longest Leg First + Familiarity</td></tr>
<tr><td rowspan="6">Evaluation Metrics</td><td colspan="5">Route Length</td></tr>
<tr><td colspan="5">Number of Turns</td></tr>
<tr><td rowspan="4"></td><td colspan="2">Familiarity</td><td colspan="2">Errors</td></tr>
<tr><td rowspan="3"></td><td rowspan="3"></td><td rowspan="3"></td><td>User Route Length</td></tr>
<tr><td>User Route Number of Turns</td></tr>
<tr><td>User Route Similarity</td></tr>
</table>

Table 3.1: Phase Additions and Modifications.

To provide continuity, only a small number of metrics are examined at each stage but these are carried through multiple phases. Combining these with the reuse of environments and test conditions wherever possible, gives credible comparisons throughout the project. The metrics used are defined in some detail below:

- Route Length – Length of route given as a percentage of the length of the route produced by the Shortest Path heuristic for the same start and end points.
- Number of Turns – Number of turns in route given as a percentage of the number of turns in the route produced by the Minimum Turns heuristic for the same start and end points.
- Familiarity – A measure of the combined familiarity of all components in a route. This is shown for all routes produced by the appropriate heuristics.
- Errors – Number of errors occurring along each route, given as a percentage of the maximum

possible errors (assuming one error per component within the route). The box plots for this metric show the 5th, 25th, 50th, 75th and 95th percentiles.

- User Route Length – Length of all routes selected by the system, given as a percentage of the length of the route supplied by the user for the same start and end points.
- User Route Number of Turns – Number of turns in all routes selected by the system, given as a percentage of the number of turns in the route supplied by the user for the same start and end points.
- User Route Similarity – For all routes selected by the system, the percentage of components within a route which correspond to components within the route supplied by the user for the same start and end points.

The following sections and chapters show how the phases were constructed and evaluated in order to create the final application. From the environment representation to the user simulation, they detail all of the processes involved and how each system element was built.

3.2 Phase 1 – Building the Grid and Creating Simple Routes

One of the most important steps in constructing a route creation system is also the first. Before any route heuristics or wayfinding techniques can be considered, a representation for the environment must be defined. This section will describe the how initial stages of implementation were conducted, and how simple representations of artificial surroundings were built. It will then go on to explain that by applying basic route selection heuristics to this, elementary routes were formed.

3.2.1 The Grid

	Complexity Level		
Route Components	Initial	Sufficient	Complete
Path	x	x	x
Node	x	x	x
Landmark		x	x
Edge			x
District			x

Table 3.2: Enviroment components by complexity level.

The approach to be used to form the representation is taken directly from Lynch [26] and can be divided by complexity as shown in Table 3.2. Phase 1 of development corresponds to the initial level of complexity and all of the components within it. In order to simplify the representation, each path is broken down into links (sections of path occurring between two nodes) creating unique components to be traversed.

The system map is designed to be input to all areas of the code, and used for both route creation by the main algorithm, and route walking by the user simulator. It contains only two lists:

- Route components – instances of all the nodes and links within the environment.
- Routes – a pointer to each node and path within a pre-loaded route (these routes are intended to represent information known to the user prior to the initial system use).

The data contained within these components and routes is shown in Table 3.3.

Component	Link	Node	Route
Data	Link ID Link Name Connected Nodes (Pointers)	Node ID Node Name Connected Link Details	Route ID Route Length Route Components
Container	Link Details		
Data	Link (Pointer) Angle		

Table 3.3: Phase 1 component data structures.

By constructing the components in this way, much flexibility is achieved within a relatively simplistic structure. As links contain no information about actual locations, they may be straight or contain many curves without modifications being required, meaning that the environment can range from a straightforward grid to a realistic map and anything inbetween. This also applies to nodes, with the end result being a map that can contain any number of components corresponding to either real or stylistic environments. To avoid issues with graphics and user interfaces, the grid was input to the system through a text file, using a suitable format for each component.

Although a grid is a very artificial environment, it has several advantages for both creation and testing. Grids are very easy to build automatically, making it possible to construct and test several different maps quickly without encountering the issues associated with converting real environments into a suitable form. Also, as each link can be defined as having unit length, suitable cost variables and values can be found with only minimal experimentation.

3.2.2 Route Creation

As the aim of this project is to produce a cognitive route, this algorithm must create more than a shortest path solution and this is achieved in two ways:

1. The use of known human route heuristics when in an unknown environment or familiarity data is unavailable.

Begin with the start node and repeat until the end node is found: • Travel to each neighbouring node and calculate the cost of this journey. • Mark the current node as fully explored. • Find the minimum cost node which is not fully explored to process next. Find the minimum cost to travel from the start node to the end node. • Continue the process until all routes are at least this long.

Table 3.4: Dijkstra's algorithm

2. Extending these heuristics and adding to them with routines involving user knowledge as it is gathered and becomes available.

By using these techniques, multiple routes can be formed and suggested by the algorithm, and then passed to an evaluation stage which decides on the most appropriate.

Despite their differences, most of the heuristics to be considered are implemented using variations of Dijkstra's algorithm [10], with modifications to the route cost in each case. This approach is straightforward and simple to implement, and has been proved to be the most appropriate method for this type of wayfinding by Zhan [41].

As the whole grid is input and therefore available, the approach used was one of travelling through a familiar environment from and to at least partially known locations. To achieve this, three of the most commonly used basic heuristics [13] were chosen for this phase:

1. Shortest Path – Minimises route cost with respect to its total length.

2. Minimum Turns – Minimises route cost with respect to the number of turns within it.

3. Longest Leg First – Forms the longest chain of links before the first turn as is feasible.

The first two employ only Dijkstra's algorithm as given by Table 3.4 to find the minimum cost route, with variations to how the cost is calculated. For the Shortest Path heuristic, path length alone is used as the cost, whereas for the Minimum Turns approach an additional cost for each turn is applied.

The Longest Leg First heuristic is somewhat more complicated, requiring further processing of the components. In this context it uses the same procedure as the Minimum Turns, recording not only the least cost path, but also a slightly higher cost alternative which has a longer distance to its first turn wherever possible. Given both a best move and a slightly worse one, it compares the distance to the first turn in each instance. If the alternative step has a longer distance, and hasn't been encountered before, this is selected over the least cost move. By performing this comparison, the route with the longest distance before the first turn can be found, and once complete, any points at which the path doubles back on itself are removed.

3.2.3 Evaluation

3.2.3.1 Method

For this initial phase of development, a simple grid containing only nodes and links was constructed and from the 100 nodes included, 20 pairs were selected to represent the route start and destination points required for the route creation task. By randomly removing a number of paths (between 15 and 40 in this case), but being careful to avoid disconnecting any of the selected node pairs, a further ten deformed grids were created. Figure 3.2 shows one of the resulting representations, and the marked nodes indicate the pairs selected for testing.

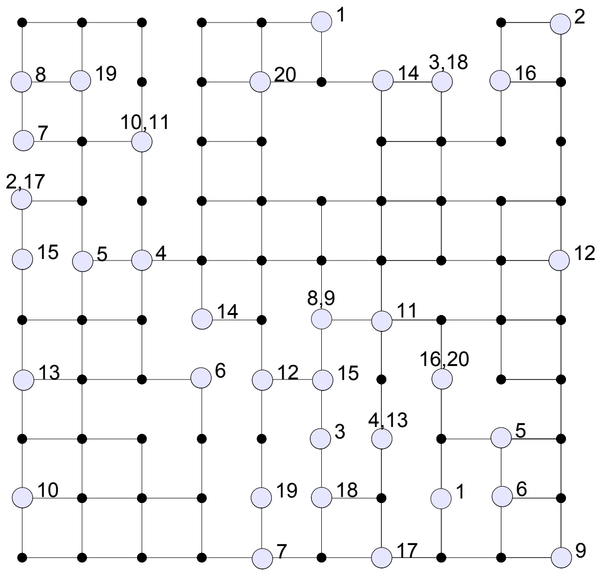

Figure 3.2: Basic Environment: Deformed grid environment representation

The heuristics all followed the same basic procedure to create a route, with variations to cost or node processing as described previously. In addition, a number of supplimentary rules were followed:

1. As each node is discovered, record the node immediately prior to it and the cost of the partial route to this point.

2. If a node is rediscovered by an alternative partial route, only replace the recorded route details if the current route has lower cost than the original.

3. Once all partial or complete routes of a suitable size have been generated, begin with the end node and work backwards chaining through the connected recorded previous nodes, noting each one, until the start point is reached.

To evaluate these heuristics, they were used to create possible routes between each selected pair of nodes on each map. By taking this approach, 600 comparable solutions (10 grids * 20 node pairs * 3 heuristics) were produced which enabled several route attributes to be examined.

3.2.3.2 Results

One of the simplest tests which can be used to evaluate suitability is to compare the length of each of the routes produced. Obviously the Shortest Path heuristic will always give the shortest route between two points, but the performance of the other approaches can easily be judged against this. Figure 3.3(a) illustrates the relationship between the results of each heuristic when used to create the routes. It indicates that in all but a few routes, the lengths of the solutions produced by all of the heuristics are equal. The routes which do display some variation account for 18%, but they do show that the performance of the Minimum Turns approach is worse than the other two in a small number of cases.

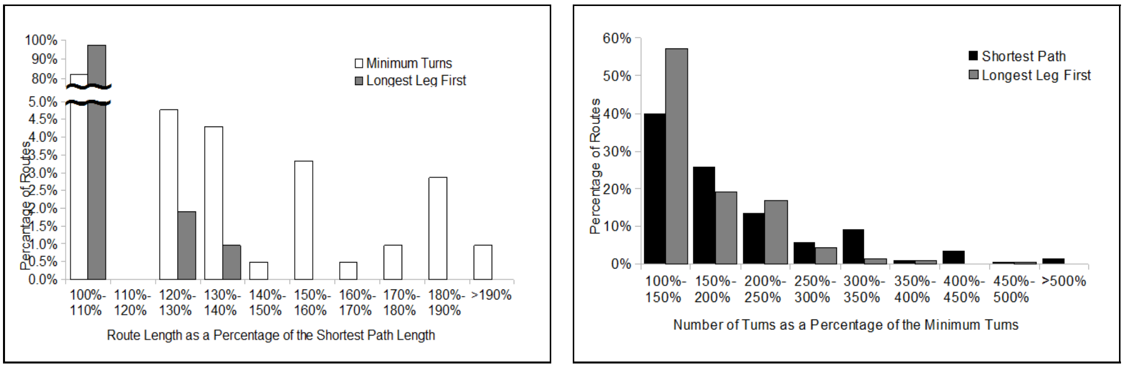

(a) Route Length as a Percentage of Shortest Route Length (b) Number of Turns as a Percentage of Minimum Turns

Figure 3.3: Basic Environment: Route Metrics

A second evaluation technique is that of comparing route complexity, creating a more cognitive metric. This complexity can come from a variety of route attributes including its structure and the prominence of its components [3]. One of the most commonly used elements with respect to route choice is the number of turns required to travel between start and destination points. Using this metric, a higher number of turns corresponds to a greater complexity and therefore less desirable route. Figure 3.3 compares this measure of complexity in the routes created by each heuristic. Unlike the previous metric, it indicates a significant difference in the performance of the individual approaches, with the Shortest Path heuristic creating the most complex, and therefore worst, routes. This variation occurs in 60% of the Shortest Path solutions, but also in approximately 43% of the Longest Leg First routes.

3.2.3.3 Discussion

Despite producing the longest routes, the second metric indicates that the Minimum Turns heuristic gives the best overall results for this phase. The problems with route length created by this approach could be resolved simply, by applying a threshold to this attribute when calculating cost during the route creation process. Using this heuristic constitutes a more cognitive approach than that of the Shortest Path, but considers only one aspect of human wayfinding techniques. A better solution may be found by combining this with other appropriate principles, and this will be investigated in the following sections.

Chapter 4

The Virtual User

This chapter introduces the concept of a virtual user to the system, allowing more human wayfinding principles to be integrated into the process. It covers phases 2, 3 and 4 of the development, concerning the user memory, landmarks and user simulation, and represents the psychological side of wayfinding. From Figure 3.1 it can be seen that these additions will complete the system, and the following sections will show how the complexity of the functionality is gradually increased to reach this point.

The first section in this chapter will discuss how a user memory was implemented, allowing familiarity to play part in route creation and increasing the attributes available to the heuristics. Following this, the environment model will be extended to encompass landmarks, bringing this up to the sufficient level (see Table 3.2). Finally, section 4.3 will explore the process employed to build a user simulator, created to automatically evaluate the routes produced and permitting the 'best' solution to be found. This is the last layer of intricacy and, as shown in Table 3.1, provides additional metrics to be examined.

4.1 Phase 2 – Implementing User Memory

Familiarity is an important factor affecting both the creation and successful navigation of routes by most people [26, 24]. When examining a map of a partially experienced environment to plot a route, most individuals will first locate known areas and use these to form relationships between known and unknown destinations. To approximate this behaviour, the user memory is designed to be a repository for information on routes experienced by the individual concerned. Route components encountered during traversal are stored within it, and used to create solutions and simulate wayfinding in all areas of the system. In reality, it models the most basic form of cognitive map, storing wayfinding data in a way similar to that found in human subjects.

User Memory	User Component
User Map $-$ User Components	Route component pointer
User Route $-$ component IDs	Route IDs
Loaded flag	Familiarity
	Just seen flag

Table 4.1: Phase 2 user memory structures.

4.1.1 Component Storage and Decay

The storage of component details and the decay of information due to forgetting form the two processes involved in user memory. Combining these functions gives the same effect as trail laying in swarm intelligence [4].

During the storage element, for each component passed to the user memory the algorithm is first checks to see if it is already known. If it is, then the familiarity is simply increased; if not then more processing is required. When an unknown component is received, it is initially stored as a direct copy of the incoming data, along with which route it was encountered on. Once the current run is completed, the whole map is checked and some information modified. This is done to retain details of only known or important components rather than all available data, and is performed in the following way,

- Check each attached component (links or nodes),
  - if they are not also held in memory then remove this attachment.

By retaining the number of links attached to a node, but removing unexplored details, this method preserves the structure whilst limiting the data stored. In both the known and unknown cases the ‘just seen’ flag is set to indicate that it has been seen during this run and should not be affected by the component decay stage.

At the end of each run the memory is updated and executes the component decay functionality to any parts of the map without the ‘just seen’ flag set. The update iterates through each component and checks its familiarity value. If this is above zero, it is multiplied by the decay value and converted back to an integer (to prevent leveling off below one). However, if the value of familiarity is already zero, data within the component begins to be forgotten in a specific order according to how many updates have occurred since it was last seen:

1. If the connected component pointers exist, replace these with null values.
2. If associated landmark pointers exist, replace these with null values.
3. Remove the component from memory.

4.1.2 Using Memory for Routes

Integration of memory within route creation could be implemented in several different ways, from considering just whether or not a component is known, to how frequently and recently it has been visited. In this case, a cost associated with the actual familiarity stored within the component is added to each of the basic heuristics from phase 1 using the following equation:

$$Cost = (FamiliarityConstant * familiarity) + OriginalCost$$

The most suitable level for the familiarity constant was found to be 0.1 through a large amount of testing during this phase, although the results are not given here. Too large a value will cause the route to use the most familiar components, but lose any connection to the basic heuristic. If the value is too small, any effects associated with the familiarity will be overwhelmed. These modifications increase the number of approaches available for route creation to six, with the new ones being as follows:

1. Shortest Path + Familiarity – Minimises route cost with respect to its total length, but maximises it with respect to familiarity.

2. Minimum Turns – Minimises route cost with respect to the number of turns within it, but maximises it with respect to familiarity.

3. Longest Leg First – Forms the longest chain of links before the first turn as is feasible, whilst incorporating as many known components as possible

4.1.3 Evaluation

4.1.3.1 Method

The experimentation in this phase uses the grids produced previously, but to keep the number of results manageable, only the first five were selected for testing during this phase. For consistency, the original 20 node pairs were again utilised as the start and destination points, but this time applied to the task in sequences rather than singularly. This was done by forming a list of the point pairs, and then producing a route for each pair in turn. The solutions suggested are stored in memory as they are created, providing increasing knowledge which will then be available for future use. To give a fair comparison, the process was then repeated with the start point in the list incremented, and continued until all points have begun the operation. This approach allows each pair to occur at all positions in the cycle, and they are therefore tested with a variety of existing familiarity levels. A single heuristic was used for the route creation process, with the procedure being completed in its entirety for each of those available, resulting in the individual criteria being tested independently.

Initially, the system was tested with perfect memory with a decay multiplier set to one allowing the full effects of familliarity to be examined. Further experiments were performed with the decay rate raised to two, meaning that each component held in memory but not visited by the most recent route has its familiarity value halved. In all tests the familiarity constant was set to 0.1, the value found to give the best balance between costs earlier.

4.1.3.2 Results

The results discussed here apply to only routes where the user memory has some content, enabling a true comparison of both basic and memory heuristics across all metrics. This indicates that at least one route must have been successfully generated and stored before the new heuristics were tested.

When integrating a factor such as familiarity, the immediate question is how will this affect the length of the routes produced? Figure 4.1 illustrate the impact of this modification for both perfect and decaying memory. From Figure 4.1(a) it is clear that an increase in familiarity will produce a corresponding increase in length for a small number of routes. This trend would be expected to continue over larger number of traversals, and the number of solutions with increased length growing significantly. By allowing the memory to decay, and therefore sacrifice some familiarity, this trend can be halted. Figure 4.1(b) shows that with a decay of 0.5 there is no increase in the length of routes produced, even with component knowledge still playing an important part in the creation process.

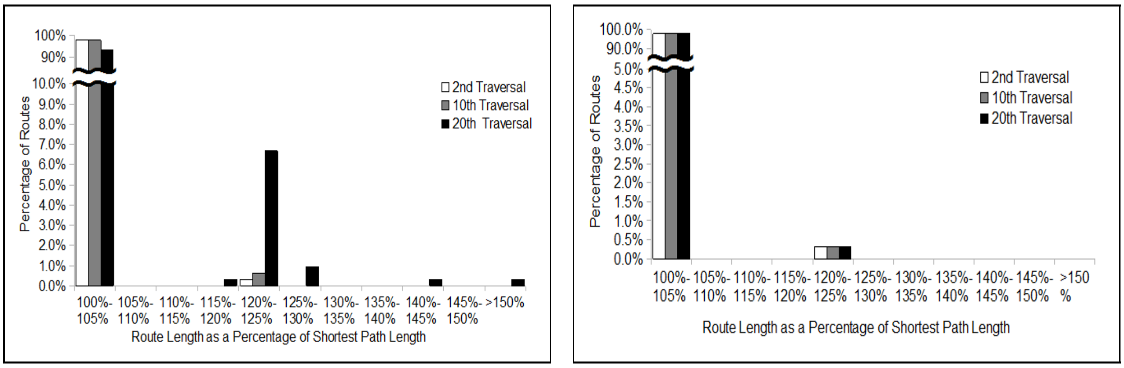

(a) Perfect Memory (b) Decay of 0.5

Figure 4.1: User Memory: Route Length as a Percentage of Shortest Route Length as Familiarity Increases

The actual levels of familiarity within the routes are shown in Figure 4.2. Despite having no direct influence on the cost related to construction, known components are likely to exist within routes created by the basic (phase 1) heuristics. The familiarity associated with these for both perfect and decaying memory are given in Figure 4.2(a) and Figure 4.2(c) respectively, and may be considered to be the baselines. By examining these, it is obvious that allowing the memory to deteriorate over time has a serious effect on the familiarity of routes. Although some reduction was expected, it is much larger than anticipated. Figure 4.2(b) and Figure 4.2(d) however, show that in both cases the phase 2 memory heuristics increase these levels, indicating that they are working correctly.

Figure 4.3(a) shows how the increase in route length breaks down across the heuristics, mostly

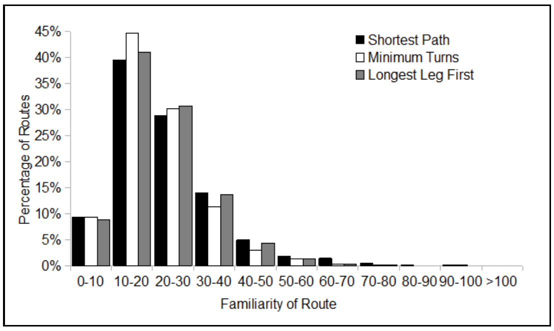

(a) Basic Heuristics (Phase 1) – Perfect Memory

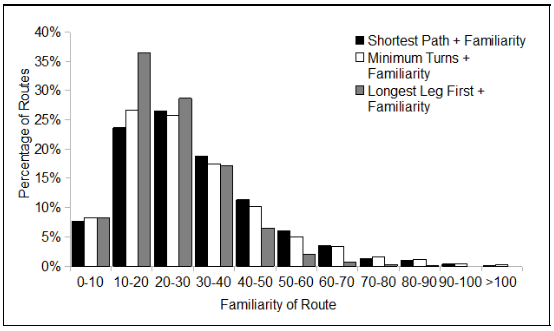

(b) Memory Heuristics (Phase 2) – Perfect Memory

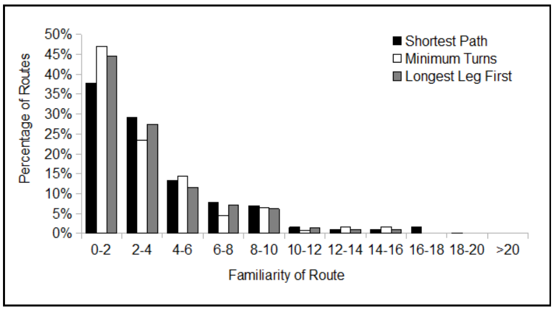

(c) Basic Heuristics (Phase 1) – Decay of 0.5

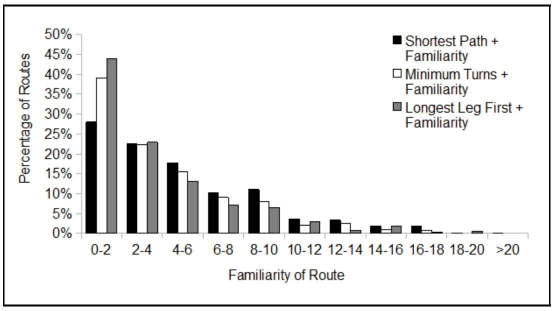

(d) Memory Heuristics (Phase 2) – Decay of 0.5

Figure 4.2: User Memory: Route Familiarity

due to the Minimum Turns + Familiarity approach. From Figure 4.3(b) it can be seen that these increases are removed by the decay functionality, with the Minimum Turns + Familiarity method now performing even better than its phase 1 counterpart. The reason for this becomes evident when Figure 4.4(a) and Figure 4.4(b) are compared. They indicate that with perfect memory, the Minimum Turns element of the heuristic becomes overwhelmed by the familiarity, performing worse than the Longest Leg First alternative. By forcing the loss of stored data by decay, this influence is removed and the heuristics efficiency returns to its previous levels.

One of the most suprising results in Figure 4.4 is that, in a small number of routes, the phase 1 Minimum Turns heuristic does not produce the solution with the lowest number of turns (indicated by the values of below 100%). This is due to the approach taken for creating these routes, and is an anomaly rather than an error in the code.

4.1.3.3 Discussion

During this phase, only a relatively low number of traversals are performed during each iteration, and the increases in route length would be expected to magnify as this rises. This effect would seem to rule out the use of a perfect memory, but as a tradeoff between familiarity, length and complexity will always be present, the optimum decay value is likely to occur somewhere between the two levels tested. In addition it should be noted that the familiarity constant was set with a decay value of 0.5, and adjusting this may help the situation. Investigation of these hypotheses was however beyond the

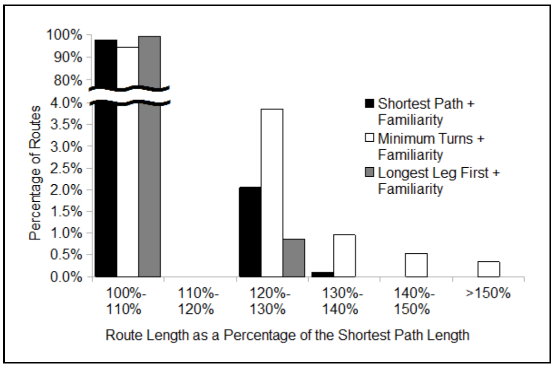

(a) Perfect Memory

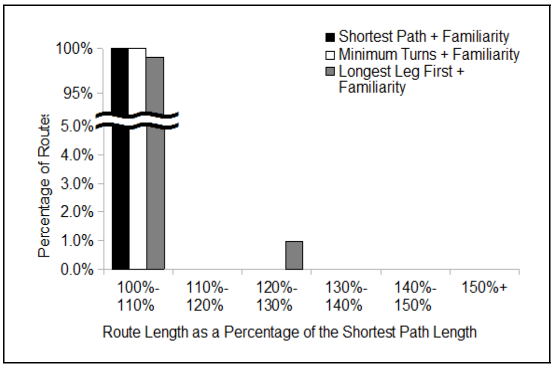

(b) Decay of 0.5

Figure 4.3: User Memory: Route Length as a Percentage of Shortest Route Length

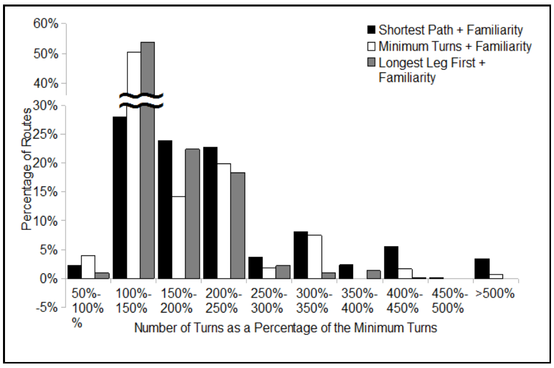

(a) Perfect Memory

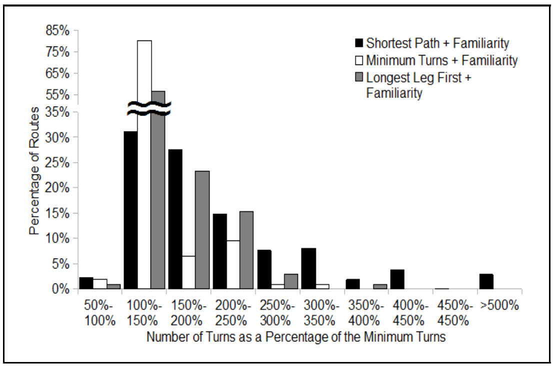

(b) Decay of 0.5

Figure 4.4: User Memory: Number of Turns as a Percentage of Minimum Turns

scope of this project.

The unexpected results pertaining to number of turns can be explained by examining the implementation of the Minimum Turns heuristic. In this context, the algorithm employs a local optimisation technique, meaning that for each node explored, only the best partial route to that point is recorded. The cost involved in calculating the preferred course of action is gathered from this partial route alone, with no consideration of the remaining cost required to reach the destination point. Although this method does not produce perfect results every time, it is faster and requires smaller processing and storage capabilities than global alternatives.

4.2 Phase 3 – Adding Landmarks

As a successful route will need to include landmarks, it was important to incorporate these into the model during this phase, increasing the representation complexity to the sufficient level. In addition, landmarks should also have some impact on the routes created by the wayfinding heuristics, and solutions for these two issues will both be considered in this section.

4.2.1 Extending the Environment

Expanding the environment to accommodate landmarks means first defining the important attributes needed to be held within the structure. Unlike nodes and links, very little data about landmarks is actually required to form routes. To reflect this, only a basic component was created containing the landmark's ID and name, along with a variable signifying its strength. This strength represents a single value found by combining measures of the prominence, singularity, prototypicality and content, and is an indication of how likely the landmark is to be recognised.

Component	Link	Node	Landmark	Route
Data	Link ID Link Name Connected Nodes (Pointers) Landmark Details	Node ID Node Name Connected Link Details Landmark Details	Landmark ID Landmark Name Strength	Route ID Route Length Route Components
Container	Link Details		Landmark Details	
Data	Link (Pointer) Angle		Landmark (Pointer) Distance Angle	

Table 4.2: Phase 3 component data structures.

Adding landmarks to the environment involves more than just including the new components in the system map. To make them useful, they must also be incorporated into the existing components in a pertinent way. Table 4.2 shows how this was achieved by attaching the landmarks to both adjacent links and nodes (more details of this will be given below), along with specifying the angle at which they can be found. For completeness, it is also important to give the distance of the landmark from the component to which it is connected, as this can be used to consider the true strength of the reference point.

Finally, consideration must be given as to how landmarks encountered by the user can be stored in memory. To keep this process as simple as possible, a straightforward approach to recording the required information was taken. As each node or link on a route is traversed, it is assumed that all landmarks connected to this component are recognised and therefore stored. Once in memory, this type of feature is treated in the same way as the other components, with familiarity incrementing and decaying as described in section 4.1.1.

4.2.2 Evaluation

4.2.2.1 Method

To allow for direct comparison with the results for phase 2, the grids used for testing in that stage were modified to include a number of landmarks. Up to 71 of these were placed randomly around each grid, positioned centrally to the cells formed by the nodes and links. These were then connected to the surrounding components, forming a sufficient environment. Figure 4.5 shows the same deformed grid as was shown in Figure 3.2, but with its automatically generated landmarks marked. The node pairs shown on the earlier grid will be reused but are not illustrated here.

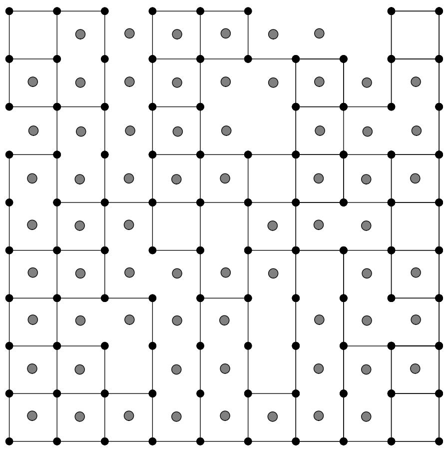

Figure 4.5: Landmarks: Deformed grid environment representation

Other than the grids, an identical approach to testing was adopted to that of phase 2, with the same start and end points being used. All variable values were kept constant, but only memory with a decay of 0.5 was investigated at this stage. The only modification made to the algorithm was the inclusion of landmark familiarity in the cost calculations where appropriate.

4.2.2.2 Results

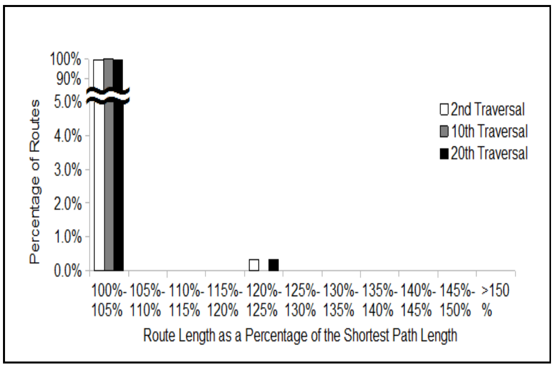

Figure 4.6: Landmarks: Route Length as a Percentage of Shortest Route Length as Familiarity Increases

As with the decaying memory in phase 2, Figure 4.6 indicates that increasing the familiarity of the route has little or no effect on its length. This is a highly desirable outcome, especially when the results shown in Figure 4.7 are examined more closely. For both basic (Figure 4.7(a)) and memory (Figure 4.7(b)) heuristics, the levels of route familiarity show small but significant increase across all approaches, indicating that more known components are being used. The biggest rise is found in the Shortest Path heuristics, with the familiarity variation outperforming all of the other methods.

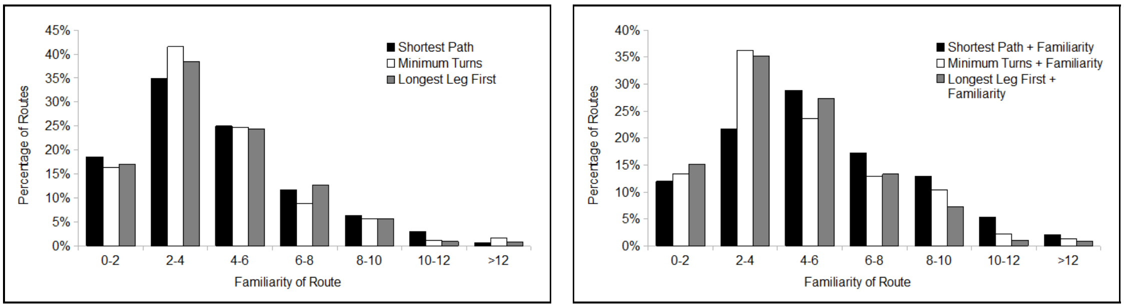

(a) Basic Heuristics (Phase 1) (b) Memory Heuristics (Phase 2)

Figure 4.7: Landmarks: Route Familiarity

Figure 4.8(a) shows a near identical series of results for the route length when compared to those in phase 2, with only a slight increase in the performance of the Longest Leg First + Familiarity heuristic. As illustrated in Figure 4.8(b), only small differences are also seen when examining the number of turns metric. Here the Shortest Path + Familiarity seems to have the largest gain in suitability.

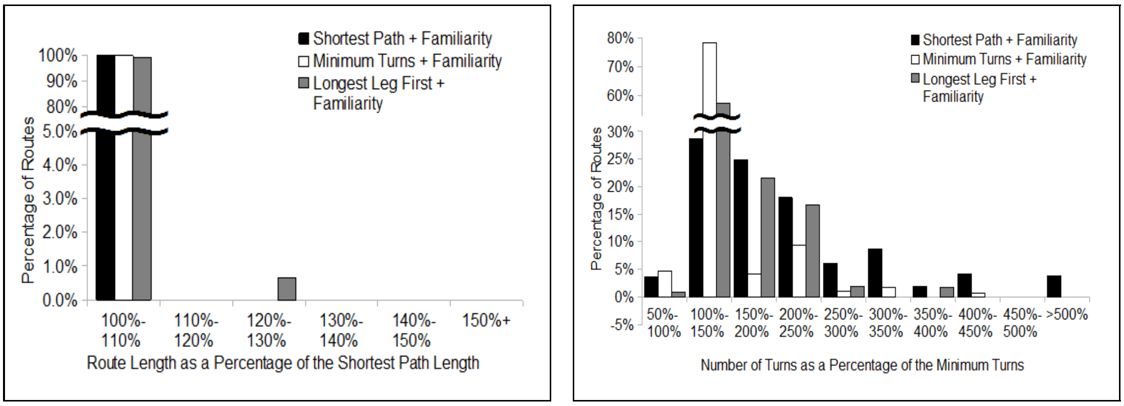

(a) Route Length as a Percentage of Shortest Route Length (b) Number of Turns as a Percentage of Minimum Turns

Figure 4.8: Landmarks: Route Metrics

4.2.2.3 Discussion

The addition of landmarks to the environment increases the familiarity of the routes created without extending their length. This is a very useful observation, and gives justification for the use of an environment with a sufficient level of complexity in this situation. In most cases landmarks can be seen from a number of nodes and links, meaning that their familiarity can be increased by different routes which have no intersections or common components of any other type. This is a unique attribute amongst the components of this level of environment, making them particularly applicable to the task at hand.

4.3 Phase 4 – The User Simulator

User simulation provides an automatic way of evaluating the validity and robustness of a route without human intervention, by mimicking the behaviour expected of a real traveller. It also allows the stability of the solution to be tested against individual components of human error, and variation of each of these elements without the need for several test subjects. This phase uses typical wayfinding errors as a post-creation evaluation step, examining the the different routes constructed to select the best.

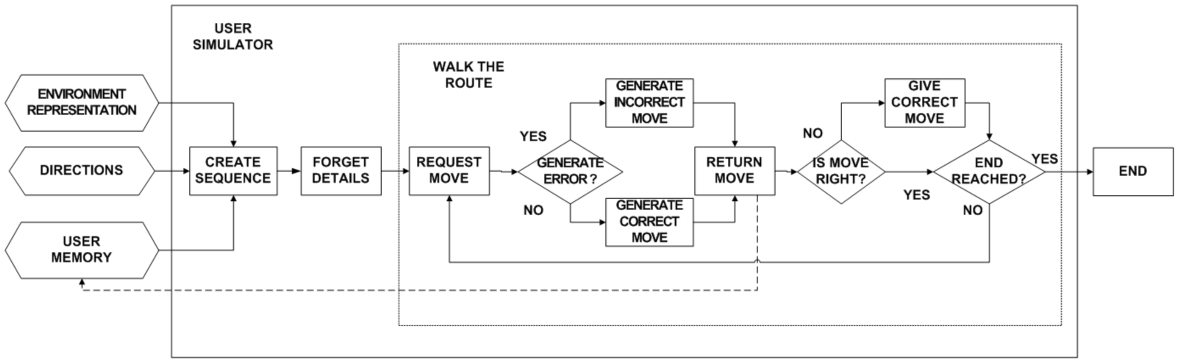

Figure 4.9: User Simulation: Data Flow

The user simulator is designed to approximate the behaviour of a human in a real-world experiment. To achieve this it follows the same series of steps as would be expected in a wayfinding experiment using actual individuals (as shown in Figure 4.9, taking the map and directions, processing them into a series of moves, and then traversing the route. In this case there are three separate inputs:

- Environment representation - model of the area to be traversed.
- User memory - a priori user knowledge of the routes and environment involved in the task.
- Directions - transfer of information about the route to be followed from the tester to the user.

From these a sequence of moves is formed, and the effects of data loss during this process simulated. Once the sequence processing is complete, the algorithm simulates walking the route by iterating through the following three steps until the end of the route is reached:

1. The system requests the next move from the user.
2. The user either provides a move from the information provided by the sequence, or an incorrect move generated by the error functionality.
3. The system checks the move provided by the user against the next move in the actual route. If the two aren't the same it corrects the user, giving it the actual move from the route, and records details of the error.

In addition, if valid, both correct and incorrect moves can be stored in the user memory if required.

4.3.1 Producing Directions

Directions are one of the main ways of transferring information about a route [27], in this case to either a virtual or human user. To be useful, they must contain a series of locations, relevant landmarks (where available) and actions. In order to achieve a solution which successfully completes this task, the format of the directions must be carefully considered. They need to be both readable by humans, and contain enough information to be applicable to the user simulation. The simplest way to accomplish this is to convert all the relevant data for each step along this route into sentences representing the location, action and landmark concerned. It is important that enough information is provided to enable the user to successfully navigate between the start and destination points without including unnecessary detail [8, 9]. Including more than one landmark may reduce the likelihood of errors being made, but increases the amount of information that has to be retained by the user, and may complicate things sufficiently to be detrimental. As there may be many landmarks available for any given node or link, the most appropriate must be found by examining strength, familiarity and position. The output of this process is a series of sentences such as:

$$At\ node\ 66\ Turn\ left\ onto\ path\ 225\ by\ landmark\ 339.$$

4.3.2 Creating the Sequence

Given a set of directions, the user simulation must convert this data inot a format that can be used to trace the route through the environment. It does this as shown below, producing up to four elements for each step.

$$At\ \underbrace{node\ 66}_{LocationID}\ \underbrace{Turn\ left}_{Action}\ onto\ \underbrace{path\ 225}_{SuppID}\ by\ \underbrace{landmark\ 339}_{LandmarkID}.$$

As indicated, the next route component is given as a supplementary ID, and used as a type of two dimensional or structural landmark. This is useful for actions where no other landmark is available.

When the sequence is complete, the algorithm applies a level of forgetting equal to the supplied percentage. This is done in a realistic manner, affecting the least familiar elements first. Initially, the location and supplementary IDs are removed, leaving only landmarks and actions. In this implementation of the user simulation, these are sufficient to traverse the route with a small number of errors. If the required level of forgetting hasn't been reached, elements in which no turn occurs are deleted from the sequence. Finally, the remaining elements are removed according to their familiarity. The resulting series of moves are then passed to the next stage of the simulation.

4.3.3 Adding User Errors

Table 4.3 shows a list of the types of errors, the likelihood for each to occur, and the method of simulation. These are all typical wayfinding problems encountered by people in both known and unknown environments [40], and are used to create a specified level of inaccuracy.

Error	Percentage	Applicability	Simulation Method
Wrong Turn	35.5%	All nodes.	If turn is left, go forward or right. If turn is right, go forward or left. If no turn, go right or left.
Missed Turn	17.2%	Turning node only.	If turn is left or right, go forward.
Selection Error	16.1%	All nodes.	Return a random link.
Misconception	10.8%	All nodes.	Return a random link.
Location Error	9.7%	All nodes.	Return a link from the same node other than the correct one.
Premature Exit	7.5%	All nodes.	Return a null pointer indicating the end of the route.
Exit Failure Error	3.2%	Last node only.	Return a link from this node.

Table 4.3: Error percentages and methods.

Not all of these errors can be made anywhere on the route, and in most cases can only occur if the current component is a node. For this phase of development, if the current position is on a link then a random node is returned, as any further processing would give a similar response. Some types of mistake are affected by the user position such as the exit failure (which can only be applied on the final node for this simulation) and missed turning error (requiring a change in direction at that point). Although returning a random link for each type of mistake associated with nodes would produce a similar effect, the methods specified permit valid moves (those where the suggested component is attached to the current location) to be recorded in user memory if required. Wherever possible, known links are used in preference to unknown alternatives. As it is possible that no mistake is actually created during this process, if no move is output the algorithm uses any known information in the same way as it would for a correct move.

How and when errors are made can also be dependent on how familiar an individual is with the current location. As alluded to earlier, object recognition is an important part of the navigation task both for familiar and unfamiliar routes. If a location cannot be recognised by any or all of its components, then an individual is far more likely to feel disorientated and therefore make an error. To replicate this effect, a combination of the familiarity of the components to the user and the strength of the landmarks involved is used to produce a recognition factor utilised to reduce the probability of an error occurring.

4.3.4 Wayfinding Assumptions

As with a human, the simulation will make various assumptions along the route if the information required is not present. Although there is very little scientific proof, it has been observed [28] that if directional data is missing at a given node, a human will carry straight on. The virtual user mimics this behaviour as long as there is some indication that this node is the one listed in the sequence. If no details of this node are given, it is assumed that the node is missing from the route and a new node is inserted with a ‘forward’ instruction. This allows for directions being given only at turning points without the intermediate points returning an error.

Additionally, all information within the sequence is compared with that available on the map for each decision. This approach will discount a single error if more details indicate the correct component. Despite this advantage, it may give contradictory moves at various points, meaning that system has to guess. In these cases, preference will be given to the IDs supplied by the sequence over any found from the action, removing issues with more than one link being possible if turning at certain nodes.

4.3.5 Evaluation

To select the most appropriate route between two points, the system must evaluate the performance of each solution against a suitable metric. This measure was chosen to be a combination of not only the number of errors occurring for each route, but also the position of the first mistake. Although the use of the first of these attributes is obvious, the latter also has advantages. The further into the route an error occurs, the more likely the user is to be able to see the destination or find it by trial and error. With these two constraints, the algorithm opts for the route which exhibits the lowest number of errors and the largest number of correct moves before the first mistake.

4.3.5.1 Method

As with the previous phases, as few modifications as possible were made to the environment and test process. To provide continuity, the grids created for phase 3 were reused, as were the start and end points. The only significant change made to the process was that the heuristics were not run independently, with all six being employed to create routes for evaluation. Once a single solution has been judged to be the best, it is stored in memory and used to generate future routes. This storage is performed with no simulated errors, recording only components in the route supplied.

Experimentation was performed with many grids and start and end points, for various levels of forgetting and inaccuracy. It was found that for values above 40% for each, the number of errors generated was larger than if an empty sequence was supplied. To avoid this and still give a suitable level of spurious moves, values of 30% inaccuracy and 30% forgetting were selected.

4.3.5.2 Results

Figure 4.10 shows how the number of errors generated for each type of route affects which is selected. It is clear that the Minimum Turns solutions produce the fewest erroneous moves in the majority of instances, and this is reflected by the number of times that its routes are selected. Although the relationship between these values is obvious, the connection in the case of the Shortest Path heuristic is much less clear. Despite producing some of the highest levels of mistakes, it is the second most selected solution. This indicates that the errors must occur much farther into the routes produced than all of the alternatives. One unexpected result is the low selection rate of the familiarity heuristics. There may be many explanations for this, but the two most likely are that there is a lack of overlap between the stored and requested routes, and the possibility that the basic heuristics are using the most familiar components, producing identical solutions to their memory counterparts.

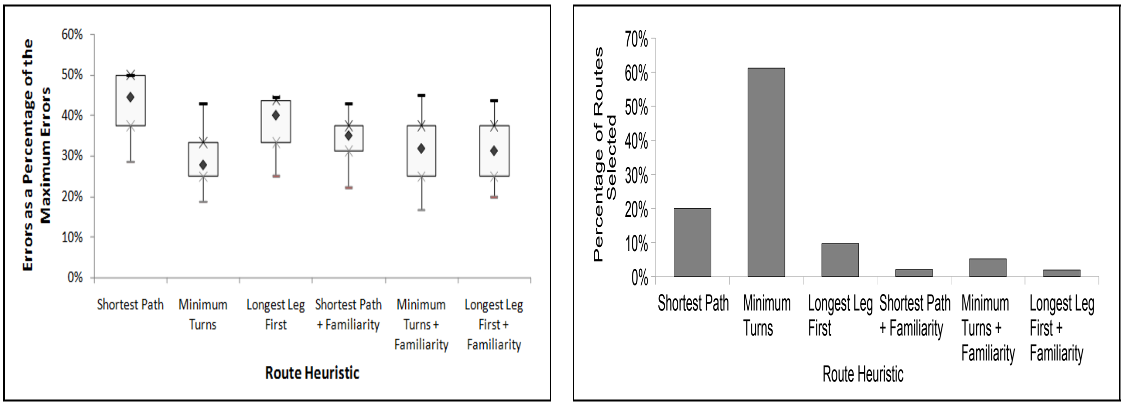

(a) Errors as a Percentage of Maximum Errors (b) Heuristic as a Percentage of the Routes Selected

Figure 4.10: User Simulation: Route Selection Metrics

By allowing the system to select any one of the six routes generated, the heuristics are no longer being tested independently. This means that there is no benefit to be gained by examining the length or complexity of all the routes created. Instead, the values of these metrics for those selected by the system to be the ‘best’ solution for each start and end point will be considered. Figure 4.11 shows these comparisons.

From Figure 4.11(a) it can be seen that over 90% of the routes selected have approximately the same length as the Shortest Path solution. In the remaining cases, the length is spread over a number of values, meaning that the errors generated do not necessarily have a direct relationship to this attribute. The same observation can be made about the route complexity displayed in Figure 4.11(b).

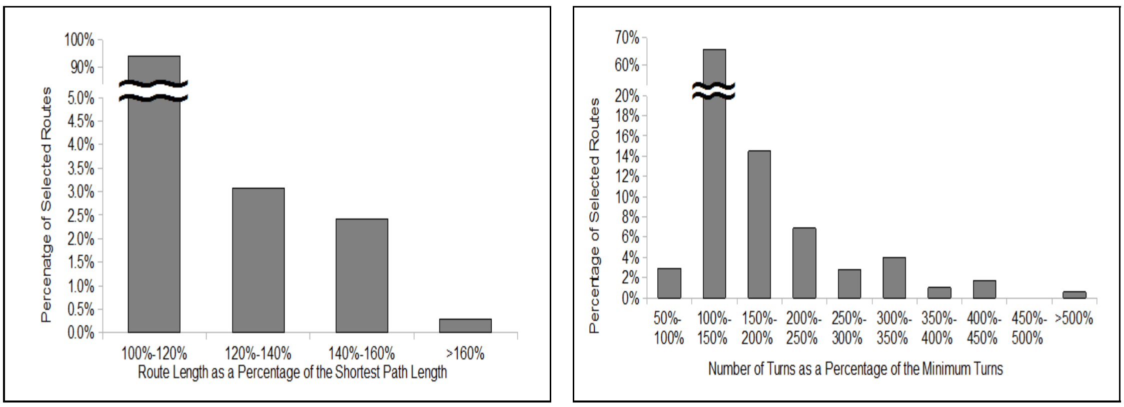

(a) Route Length as a Percentage of Shortest Route Length (b) Number of Turns as a Percentage of Minimum Number of Turns

Figure 4.11: User Simulation: Selected Route Metrics

4.3.5.3 Discussion

Although length and complexity play an important role in the suitability of a route, they are not the only considerations in selecting the most appropriate solution. By using common wayfinding errors, an alternative measure of ‘best’ can be used for evaluating the routes created. If the simulation sufficiently recreates the behaviour of a human user, this metric should be a far more reliable measure of suitability when searching for a cognitive solution. It is assumed for the purposes of this project that, despite being limited, the level of imitation produced by the algorithm is adequate for the task. This could be improved through the use of line-of-sight to future components, or memory to help replace lost information.

For the sake of simplicity, during this phase routes are recorded without error once they are selected. A more realistic approach may be to utilise the user simulation functionality for this purpose, allowing a small number of incorrect moves to be stored. Finding suitable values for the inaccuracy and forgetting required by this process, along with the suggested improvements, were beyond the scope of this project.

Chapter 5

Phase 5 – Real World Application

5.1 A Brief Introduction

Although artificial environments are ideally suited to developing and testing route creation techniques, they are not truly representative of the conditions experienced in real world settings. They have links of equal length connecting at perfect right angles, with a maximum of four links leading to every node. Regrettably, in genuine urban surroundings this exactness is rarely the case, with curved paths producing links with lengths far exceeding the Euclidean distance between their start and end points, and many of them joining a single node at a variety of different angles. These imperfections make converting the environment into a suitable model problematic, and may significantly affect the behaviour of algorithms creating routes through it.

This phase will examine how data was collected from frequent users of an urban site, and how this information was utilised to determine the most important features of the environment, and which regions are more commonly travelled. Secondly it will illustrate that a simple approach can be used to construct a deformed grid from a map of the area selected, and how the data collected can be used to influence the details held within this model. Finally, it will look at the effects of using this representation to generate routes, and how these vary from those suggested by human participants.

5.2 Constructing a Real Environment

As this project is designed to be run on real-life data, a suitable environment had to be chosen. The University of Leeds campus was selected for its convenience and variety of environments, providing a range of urban habitats in a relatively small area. A detail map of the region to be considered is given in Appendix C Figure C.1, showing features from large closely spaced buildings to open playing fields.

5.2.1 Data Collection

After defining the area to be used, the next step is to investigate how regular users traverse this environment. By collecting this information, the most commonly travelled path and well known features can be found and utilised to reduce the complexity of the representation. Diaries are a widely used and accepted method of collecting data in this field. They avoid some of the most likely issues associated with the process, whilst still providing a sufficient level of detail.

5.2.1.1 Method

To collect the required data, eight participants were selected and each provided with a diary in which to record their movements around campus during a single day. In addition, they were given an outline map showing the main buildings and pedestrian areas, and asked to draw the routes they traversed wherever possible. This two-fold approach was implemented to gain a large amount of data with minimal inconvenience to the test subjects. By combining both methods of collection, the advantages and disadvantages of each can be compared. Examples of both the diary and map used are shown in Appendix C.2.

5.2.1.2 Results

In the majority of cases a significant amount of useful information was produced, allowing the many attributes of the environment to be identified. Table 5.1 gives the start and destination points given by five of the subjects involved. The initial entry in each column is usually considered the original departure point, with each location thereafter corresponding to the end point of one route and start of the next. As can be seen, there is some overlap of locations, although it should be noted that even for identical points different routes may have given by separate individuals.

By collating the routes, a number of landmarks used for navigating around areas of the campus were found, Table 5.2. These were established by examining the text surrounding a route action for a visual reference accompanying it, these cues are assumed to be the landmarks used to locate a change in direction. They vary in size, distinction and type, and unexpectedly include districts such as old or new buildings.

Although the diary approach was successful, the route drawing was far less productive and data from this approach was discounted at this point. In addition, a small number of participants were also asked to provided directions between points which they had chosen. The following is one of the submitted passages:

> *”Union to Electrical Engineering: Go up the stairs and out the main entrance. Walk across the plaza towards the Great Hall. Cross the Worker’s court next to the great hall.*

Participant 1	**Participant 2**	**Participant 3**	**Participant 4**	**Participant 5**
MSc Study	MSc Study	MSc Study	Charles Mo	Entrance
Parkinson B11	Parkinson B11	The Edge	EC Stoner	MSc Study
Msc Study	Student Union Shop	Parkinson Building	Union Building	Parkinson B11
Roger Stevens LT8	Careers Centre	Ziff Building Accom Office	Elec Eng	MSc Study
Engineering Houldsworth LTB	MSc Study	Union Building	Union Building	Student Counselling
Equality Service	Liberty Building	EC Stoner Dol Ce Vita Cafe	Charles Mo	MSc Study
MSc Study		Roger Stevens Irene Manton Cluster		Entrance
Old Bar		Careers Centre		
Brotherton Library				
MSc Study				

Table 5.1: Start and destination points from gathered data.

Landmarks		
Edge	Alley	Shop
Car Park	Chem Block	Cromer Terrace
Library	Worther's Court	School of Psychology
Social Sciences cover	Green	Business School
Walkway	Velo Campus	Music Department
Ziff Building	EC Stoner	Roger Stevens Building
Union	Parkinson Building	Equality Service
Clarendon Place	Cafe	Geography
Whetton	School of Mathematics	George's Field
Stairs / steps	Toilets	Modern building
Red route	Gateway	Old buildings
Michael Sadler Building		

Table 5.2: Landmarks found within the gathered data.

In the back right corner there is an archway that leads to a path. Follow the path through the buildings. Keep left at Estate Services, and go up the steps to Woodhouse Lane. At the sidewalk, go left, then up the steps to the door of Elec Eng."

As can be seen the majority of these directions correspond to action-location pairs which should be perfect for comparison to the routes produced by the algorithm. Fig 5.1 illustrates the data collected by the whole process (shown in Table C.2 and Table C.3 in Appendix C) when it is overlaid on the campus map. It shows that although University of Leeds Campus is fairly large and complicated, the area covered by all the test subjects within a single day is quite small. Most of the routes invloved extend from a single, widely used thoroughfare running from the Parkinson building, past the Union, on to Clarendon Road and beyond. A smaller arm branches at the Social Sciences building, and continues down to EC Stoner.

5.2.1.3 Discussion

The method of data collection described was particularly successful and, although only a small amount of information was gathered, this was deemed sufficient in this case. An adequate number of both start and end points and landmarks were defined to assist in the task of converting the campus map. In addition, despite being difficult to trace at times, a satisfactory level of detail in routes was supplied by five of the test subjects. This allowed a quantity of suitable training and test data to be constructed.

5.2.2 Converting the Campus Map

Translating a real environment into a format suitable for this system brings its own dilemmas. With a motor vehicle, the features which can be traversed are limited to roads, car parks and similar structures. When considering pedestrian movement, those restrictions do not apply. They can travel far more freely through environments [11], using alleyways, tracks and even corridors inside buildings to reach the required destinations.

Secondly, the maps themselves may introduce their own difficulties. Even the best two and three dimensional representations rarely show the true complexity of an environment, indicating only the locations of buildings and other significant features [37]. They omit barriers such as walls and fences which can halt pedestrian motion, as well as tunnels and walkways below structures.

To overcome these issues, a set of rules should be defined and followed when constructing a grid. These should be designed to reflect the level of detail required in some areas whilst allowing for complexity to be removed where it isn't needed.

5.2.2.1 Method

To establish the traversable features, a walk of the relevant areas of campus was performed. Paths, trackways and underpasses were marked on a map, along with building entrances and other accessible features. Particular attention was paid to the areas described in the data collected, but only paths

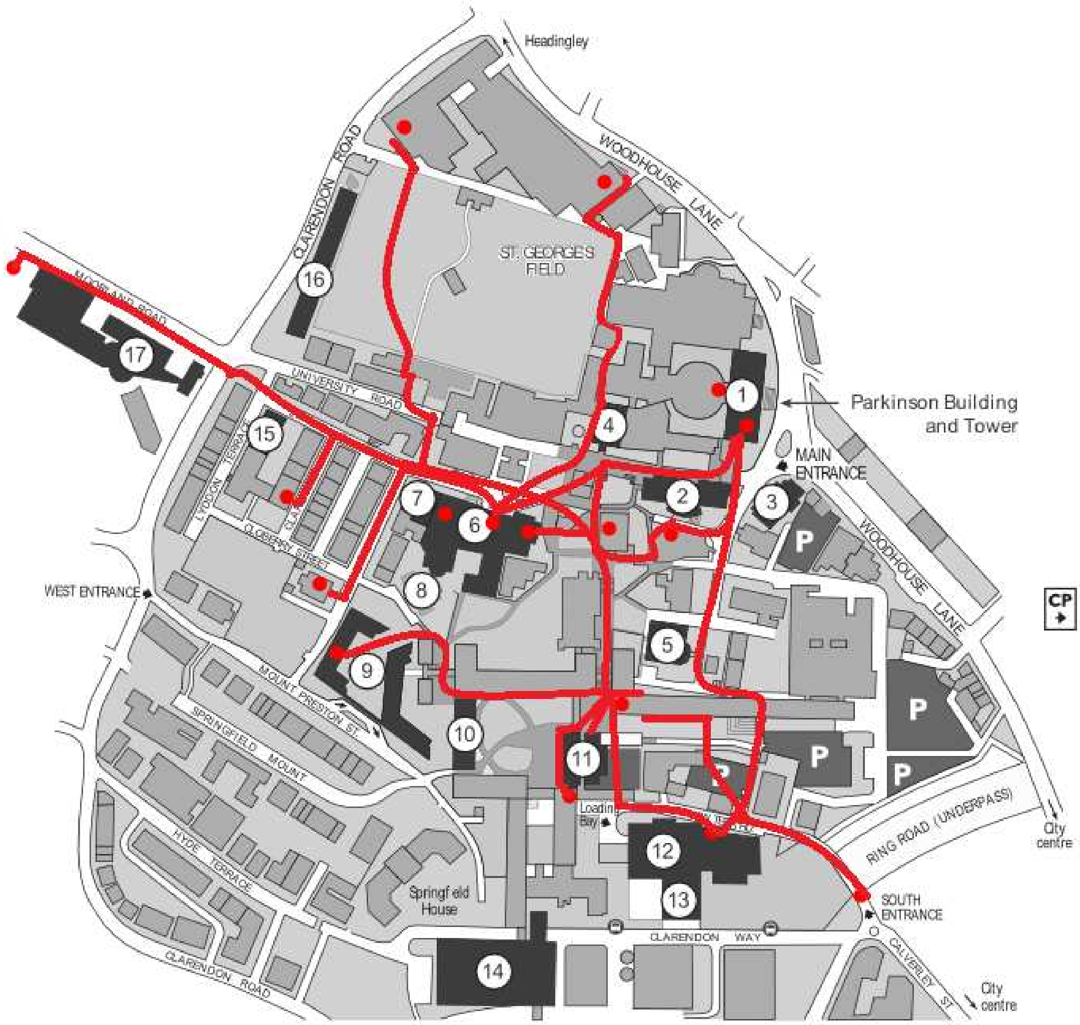

Figure 5.1: Illustration of all of the participants' routes from the collected data.

that would be considered as navigable by the casual observer. No routes which included passing through buildings or fully enclosed walkways were allowed, including the use of the 'red route'. For routes not mentioned by the test subjects, only the main roads and thoroughfares were indicated, providing simple alternatives to those given. Combining this map with the one shown in Figure 5.1, a deformed grid representing the appropriate areas of campus was formed. To be compatible with the route creation functionality, it is important that length and angle information is retained by this process, and used to augment the components created.

Although many landmarks were gathered from the user diaries, only those essential for creating sensible routes were added and restrictions on their number were necessary due to the tight time restictions placed on this project. Defining distances and strengths for this type of feature is difficult, and these values were chosen in a fairly arbitrary manner for those that were included. To provide further information to a human user, links and nodes were given names wherever possible, with multiple links having the same label if they belong to a single path.

5.2.2.2 Results

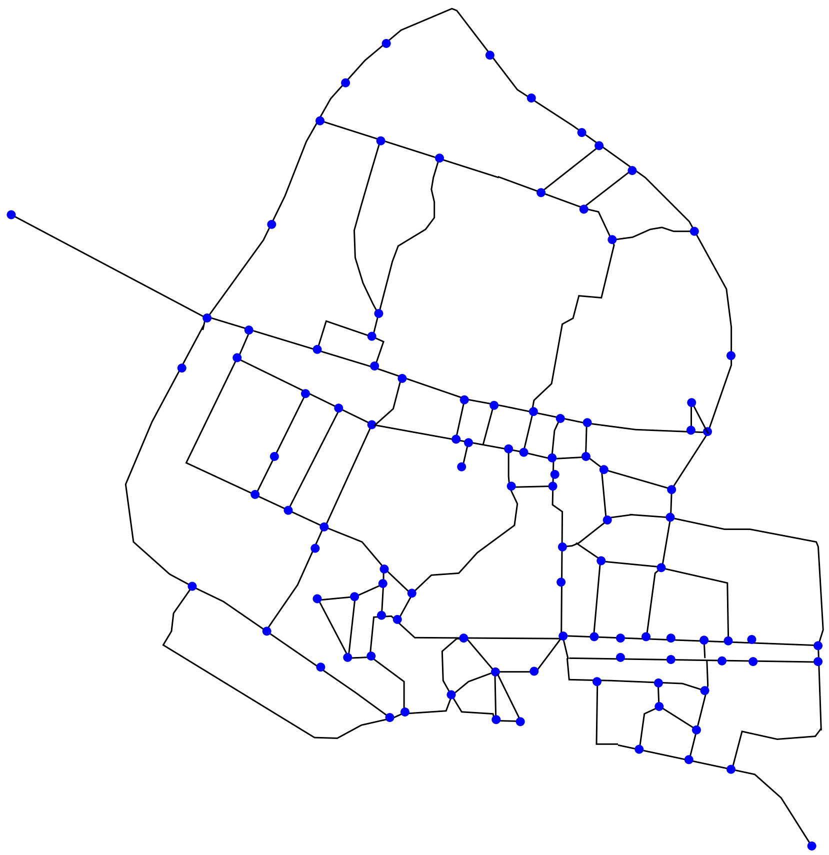

Figure 5.2: Grid created from the campus map.

Figure 5.2 shows the grid created by employing the rules and procedures described. Despite limits on complexity, 107 nodes, 133 links and 11 landmarks (not marked here) are required to describe the chosen environment. A number of the necessary nodes are used to delineate locations which are not true decision points, but indicate the entrances to buildings. These are essential for providing the start and end points corresponding to those included in the diaries. It is clear that a number of features are missing or incomplete, but the model should be sufficient for the construction and testing processes required by this project.

5.2.2.3 Discussion

This process shows that even for a restricted environment, the complexity of the grid required to fully communicate the environment increases dramatically as more details are added. Despite containing many nodes and links the model created during this phase is still very basic, giving only an adequate representation. This is sufficient in this context, but to improve the quality of the environment, many more landmarks should be included. Time restrictions imposed on this exercise meant that it was not possible to extend the model in this way.

5.3 Evaluation

As the algorithm is now complete, and a grid formed, this final phase focuses on the system's behaviour when faced with a real environment. This was an important test for the application as the previous artificial representations contain only perfect metric information, with all link lengths set to a length of one, and a maximum of four connections for each node. Modifications to the route cost functionality were anticipated but proved unnecessary, and the existing algorithm was retained in its entirety. One addition was made however, with a set of real user directions being created for each selected route. These were formed using the names of the nodes, links and landmarks, rather than their IDs.

5.3.1 Method

The grid formed in phase 5 was translated into an appropriate text file as required by the system, and for the initial stage of testing, 20 start and end point pairs were randomly selected from the nodes available. Test procedures from phase 4 were repeated, and used to check the algorithm for correct execution.

Once the validity of the system had been established further tests could be completed, with attention shifting to the user supplied data. Information from five of the collected diaries (shown in Table C.2 and Table C.3, Appendix C.3) was converted into routes, and these were used to replace the random start and end points. To create the maximum amount of training data possible, all but one of the relevant user routes were stored in memory before the route creation process began. The remaining example was used to supply the required start and end points, with the procedure being repeated for each route and each user. Solutions suggested by the algorithm were then compared with those actually used by the test subject.

5.3.2 Results

Figure 5.3(a) shows the errors generated by each route type for this real environment. Compared to the previous grids, there appears to be a far larger spread in the number of mistakes across the majority of the heuristics. Unlike the artificial environments, the best approach is the Longest Leg First + Familiarity, with the Minimum Turns routes coming in second. When combining this with the second selection criteria however, the situation changes completely as shown in Figure 5.3(b). Despite producing fewer errors, none of the routes generated by the Longest Leg First heuristic were selected as the best between the two given points. This means that these mistakes much occur far earlier than those in the solutions from the other approaches.

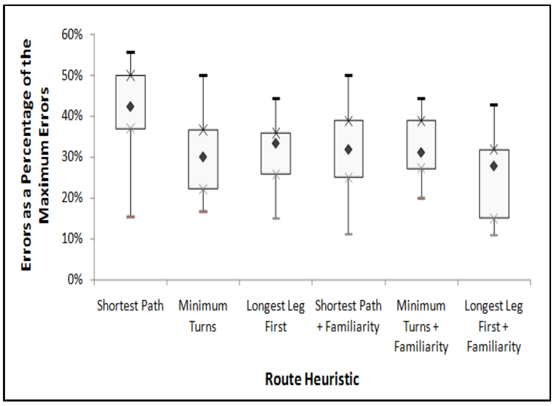

(a) Errors as a Percentage of Maximum Errors

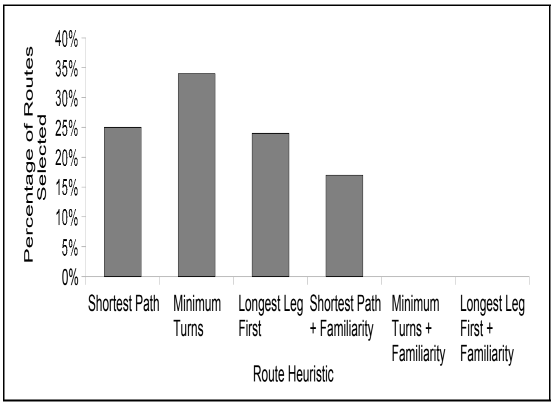

(b) Heuristic as a Percentage of the Routes Selected

Figure 5.3: Random Routes: Route Selection

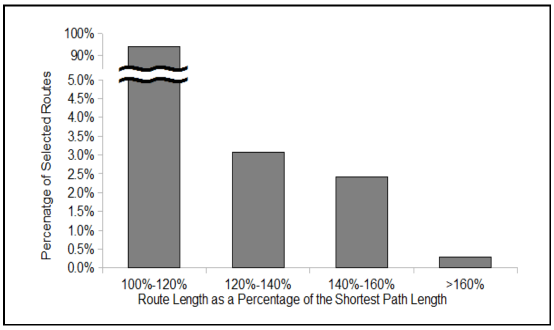

(a) Route Length as a Percentage of Shortest Route Length

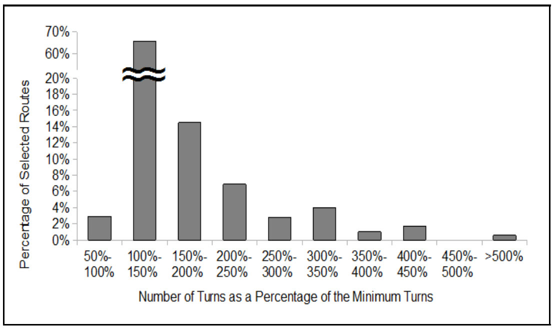

(b) Number of Turns as a Percentage of Minimum Number of Turns

Figure 5.4: Random Routes: Selected Route Metrics

The selected route metrics illustrated in Figure 5.4 are directly comparable with those in Figure 4.11, with almost no noticeable variations. Again over 90% of the routes have between 100% and 120% the length of the solutions produced by the Shortest Path heuristic. Also, almost 70% have a complexity of 150% or less, with regards to the number of turns within them, compared to those produced by the Minimum Turns heuristic.

Figure 5.5 shows how the behaviour of the system changes when it is supplied with real user data. Figure 5.5(a) indicates that the numbers of errors produced for each type of route is fairly even, with the Longest Leg First and Longest Leg First + Familiarity giving the best results. Unlike the other examples, these do translate into route selections, although not in such high numbers as expected. This can be seen in Figure 5.5(b), as can the increase in the number of routes created by the familiarity heuristics which go on to be judged as the most appropriate. At 26.67% this represents almost a 10% increase over the figures found for the random routes, and 17.55% higher than the results of phase 4.

When examining the routes suggested by the system against those stated by the user, Figure 5.6 shows that in 36.67% of cases they were identical, and in the majority of those remaining they have more in common than just the start and end nodes. This is higher than would be expected by pure

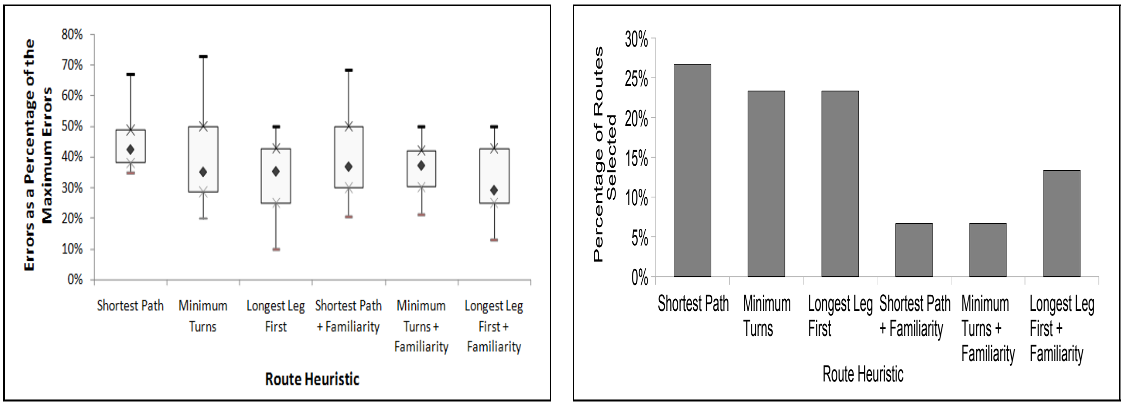

(a) Errors as a Percentage of Maximum Errors (b) Heuristic as a Percentage of the Routes Selected

Figure 5.5: User Routes: Route Selection

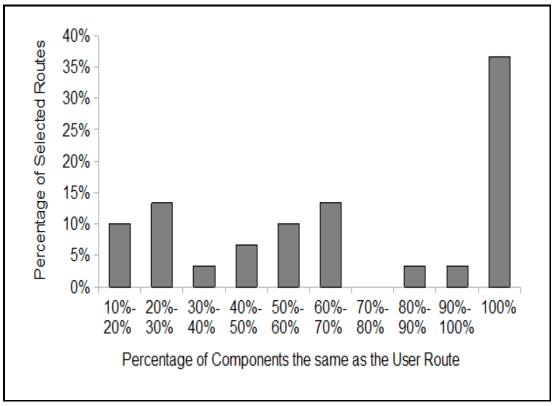

Figure 5.6: User Routes: Percentage of Components Corresponding to User Route

chance, and may indicate that the algorithm is performing well despite its limitations and small training set. Equally promising are the metrics of the routes produced, shown in Figure 5.7. All but three of the routes generated have lengths less than or equal to the user suggested alternatives (Figure 5.7(a)), and only four are more complex (Figure 5.7(b)).

5.3.3 Discussion

The first stage of testing during this phase proved that the algorithm performs as well for a real environment as it does for an artificial one. This was somewhat surprising as issues associated with varying link lengths were expected, and indicates the the route creation heuristics are robust enough to cope with this kind of change.

The user data tests produced equally successful results, with the suggested routes being valid and usable in most cases. Increases in the selection of solutions created with a familiarity based heuristic shows that a more cognitive route is found to be the most appropriate in a higher number of instances. This may be due to the amount of training data provided before any route construction began. In all

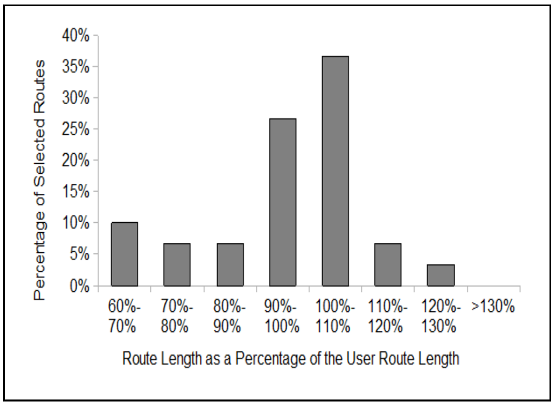

(a) Route Length as a Percentage of User Route Length

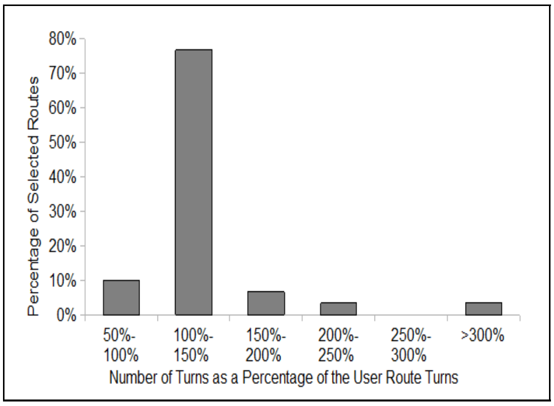

(b) Number of Turns as a Percentage of User Route Number of Turns

Figure 5.7: User Routes: Selected Route Metrics

of the previous tests the user memory was initially empty, with routes being added as they are created. This factor means that only a solution constructed with a basic phase 1 heuristic can be selected for the first route, and that several more may have to be stored before enough relevant components are present. In the current scenario, several user routes are 'pre-loaded', allowing familiarity to play a role in the creation of all possible solutions. An alternative explanation is that the overlap between routes is greater. The data collected indicates that users traversing known routes tend to prefer the use of particular route components, as suggested by Kuipers [24]. This behaviour would give a higher proportion of relevant links and nodes which have a familiarity value greater than zero.

Chapter 6

General Discussion

6.1 Other Work

Although the phases discussed so far were completed and produced some promising results, not all work attempted during this project was quite so successful. There are two instances where work was begun and then abandoned, or where the resulting output was unsatisfactory. The first of these occured during proposed work to extend the functionality of the user memory. This expansion aimed to increase the level of detail of the route knowledge stored, in an attempt to replicate more of the processes involved in the formation of a cognitive map. It was envisioned that this additional information could enable the reuse of known routes, and assist in improving the behaviour of the user simulation. Work on the user memory itself was completed, but the additional code required to integrate and utilise this resource was not started due to its complexity, and the time restrictions on this project.

A further example of unsuccessful implementation was the additional work carried out during phase five, which attempted to produce directions that were sufficient to assist a human user navigating the university campus. These directions were created for each route constructed, and were then checked to ensure that sensible routes had been selected. Most of the outputs were found to contain sufficient information to enable a frequent user to trace the route suggested, but were inadequate to allow a visitor to navigate between the required start and end points. Many of the problems with these instructions came from the lack of landmarks in the representation of the environment, but some were due to the far bigger issue of superfluous and irrelevant information. Much further time and work would be required to overcome these difficulties, and the pursuit of suitable directions was abandoned after only the initial results had been gathered.

6.2 Conclusions

This project has successfully constructed a system which uses a cognitive approach to find the most appropriate pedestrian route between two points, producing a solution which is applicable to both individuals with a loss of wayfinding skills, and those in new surroundings. It has achieved this by determining the cognitive principles to be applied to the problem, constructing a suitable model of the environment, and implementing an appropriate route production algorithm. In addition, a suitable evaluation tool in the form of a user simulator has been created, and an appropriate data collection technique identified and employed. This implementation allows several observations to be made, and a number of conclusions to be drawn.

By examining previous research, this report has shown that many factors affect human route choice, supplying criteria on the selection of pedestrian wayfinding solutions in a variety of environments. The initial phase of development assumed that a suitable environment representation can be created by using a deformed grid composed of Lynch's [26] features for a basic level of complexity. Taking this approach, a number of artificial environments were generated and used to develop and test the route creation functionality. As cognitive principles were incorporated into the system, the model was increased to represent the sufficient, but stopped short of the full level of complexity. The outcome of the process indicates that suitable routes can be produced without the need to include all defined features within the environment.

Length, complexity and familiarity have all been identified as attributes which play pivotal roles in this process, and the work here has shown that by adapting a well known route-finding algorithm to consider these characteristics, they can also be successfully employed to create suitable routes through urban surroundings. Early results showed that the traditional approach of selecting the shortest route between two points may not produce the most appropriate solution in a relatively large number of cases, and that using more cognitive criteria may increase the suitability of the routes selected.

It was illustrated that through the storage of information on environmental features experienced during wayfinding, the solutions suggested by the system can be varied automatically to consider a user's knowledge of their surroundings as it increases. The results from this phase show that this type of adaptation can have undesirable effects on both the route length and complexity, but that by allowing the stored information to deteriorate over time, these drawbacks can be limited to acceptable levels. In the following phase it was also discovered that by adding landmarks to the representation, the level of familiarity can be increased with no accompanying rise in the values of the remaining criteria.

Evaluation through user simulation has been shown to provide adequate measures of task performance by introducing human wayfinding errors and the loss of detail through forgetting. This virtual walkthrough technique allows many of the different cognitive aspects of wayfinding to be replicated, and allows a more psychological approach to be taken in determining the 'best' pedestrian route through an environment. It has indicated that although route characteristics are used to affect the success of this simulation, by combining these and other cognitive principles a more appropriate measure of suitability can emerge.

Many problems with collecting suitable data have been identified in previous studies, but the method taken by this project produced surprisingly successful results, with the majority of test subjects providing a large amount of useful information. However, several problems were experienced when attempting to produce a suitable model of the university campus, mainly due to its complexity and the time restrictions imposed on the project. Despite this, the final phase showed that the system was robust enough to produce similar results on a real environment to those found on with artificial ones, and that by comparing these to the routes gathered by the data collection process it could be shown that a number of the routes were identical. Where the two differed, the system suggested solution was comparable to if not better than that supplied, when examined with respect to the defined metrics.

In general, this project has shown that psychological principles can be applied to wayfinding tasks in a way that produces a more coginitive solution than through shortest route selection alone. Futhermore, that by simulation of human wayfinding errors and other mistakes, suggested routes can be successfully evaluated for their suitability, and that by the use of an appropriate representation of the environment, this sytem can be applied to real as well as artificial surroundings.

Future work in this area may take several different directions, from investigating specific variables within the system, to extending the functionality of existing elements, and even adding suitable visualisation modules. There are a number of thresholds and constants within the application which merit closer scrutiny. Values such as that used for decay in the user memory could be examined to determine their most effective level, or whether varying those in the route cost functions could give more robust results. The work that was begun on the storage and reuse of knowledge could be completed, as could the functionality to produce usable directions. With small modifications, attributes such as the accessibility of areas to wheelchair users could be integrated into the route creation algorithm, adapting the existing heuristics or leading to the development of new ones to incorporate this characteristic.

An alternative course may be the creation of a suitable interface for collecting data about the environment and how a human user travels through it. In addition, the selection and development of an appropriate method for presenting the resulting routes to the individual, with one approach to this being through the use of external mapping functionality such as Google Maps [19].

6.3 Project Management

As mentioned in the introduction, several significant changes had to be made to the project schedule after work had begun. Many of these were due to the number of stages of development not having been fully established before initial implementation started, but others arose from over- or under-estimation of the time required for the individual tasks. Each of these modifications were made in advance of the exercises involved, and the project plan updated as appropriate. Both the original and amended timetables are shown in Appendix D Figure D.1 and Figure D.2.

The first group of necessary changes was directly related to the type of methodology adopted during this project. The system was designed to be implemented in a number of separate phases using an iterative 'design-implement-test' approach, followed by a single period of evaluation. This procedure allowed the complexity of the application to be built up one step at a time, with different algorithms being produced for individual aspects of the functionality, and then integrated into the overall architecture. Although this was the preferred method of development for this project from its start, detailed analysis of the system elements required was not performed until after submission of the interim report, and therefore after the original schedule had been constructed. This type of modification was expected, and the initial plan contained some tasks that were vague to allow the timetable to evolve. For examples of this class of variations, please look at the activities assigned to occur directly after the user simulator milestone.

The second category of required alterations came from the miscalculation of the time estimated for each task. In the original plan, a lengthy period was allocated for algorithm design, but this was later reduced to consider just the outlining of the overall system, with the remaining design exercises being amalgamated with the relevant phase implementations. As the incremental development progressed, the complexity of adequately evaluating each separate element became apparent, leading to an increase in time allotted for this process, and a reduction in that available for further enhancements of the system. The planned extensions were therefore cut short or abandoned in favour of more productive areas of work and indepth analysis of the existing functionality. Additional time was also set aside for the final report to be written, as the initial assignment was considered insufficient.

Despite the necessity for the described alterations, all work was completed in its allotted time frame. Modifications were only made to the schedule where absolutely essential, and time was managed appropriately throughout the entire project.

Bibliography

[1] G.L. Allen. Cognitive abilities in the service of wayfinding: A functional approach. *The Professional Geographer*, 51(4):555–561, 1999.

[2] G.L. Allen. Spatial abilities, cognitive maps, and wayfinding. In *Wayfinding behavior : cognitive mapping and other spatial processes*, pages 46–80. Johns Hopkins University Press, Baltimore, 1999.

[3] Jeremy N. Bailenson, Michael S. Shum, and David H. Uttal. The initial segment strategy: A heuristic for route selection. *Memory & Cognition*, 28(2):306–318, March 2000.

[4] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *From Natural to Artificial Swarm Intelligence*. Oxford University Press, 1999.

[5] Yao-Jen Chang, Shu-Ming Peng, Tsen-Yung Wang, Shu-Fang Chen, Yan-Ru Chen, and Hung-Chi Chen. Autonomous indoor wayfinding for individuals with cognitive impairments. *Journal of NeuroEngineering and Rehabilitation*, 7(1):45, 2010.

[6] Collins. Collins dictionaries online. http://www.collinslanguage.com, 2011.

[7] Luca Latini Corazzini, Catherine Thinus-Blanc, Marie-Pascale Nesa, Giuliano Carlo Geminiani, and Patrick Pruch. Differentiated forgetting rates of spatial knowledge in humans in the absence of repeated testing. *Memory (Hove, England)*, 16(7):678–688, October 2008. PMID: 18720220.

[8] Marie-Paule Daniel, Ariane Tom, Elsa Manghi, and Michel Denis. Testing the value of route directions through navigational performance. *Spatial Cognition & Computation*, 3(4):269–289, December 2003.

[9] Michel Denis, Francesca Pazzaglia, Cesare Cornoldi, and Laura Bertolo. Spatial discourse and navigation: an analysis of route directions in the city of venice. *Applied Cognitive Psychology*, 13(2):145–174, April 1999.

[10] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, December 1959.

[11] Christian Gaisbauer and Andrew U. Frank. Wayfinding model for pedestrian navigation. In *Spatiotemporal Analysis II*, University of Girona, Spain, 2008.

[12] S Gillner, A Weis, and H Mallot. Visual homing in the absence of feature-based landmark information. *Cognition*, 109(1):105–122, October 2008.

[13] R. G. Golledge. Human wayfinding and cognitive maps. In *Wayfinding behavior : cognitive mapping and other spatial processes*, pages 5–45. Johns Hopkins University Press, Baltimore, 1999.

[14] Reginald Golledge. Path selection and route preference in human navigation: A progress report. In Andrew Frank and Werner Kuhn, editors, *Spatial Information Theory A Theoretical Basis for GIS*, volume 988 of *Lecture Notes in Computer Science*, pages 207–222. Springer Berlin / Heidelberg, 1995.

[15] Sucharita Gopal, Roberta L. Klatzky, and Terence R. Smith. Navigator: A psychologically based model of environmental learning through navigation. *Journal of Environmental Psychology*, 9(4):309–331, December 1989.

[16] Michael R. Hill. Stalking the urban pedestrian. *Environment and Behavior*, 16(5):539–550, 1984.

[17] B Hillier, A Penn, J Hanson, T Grajewski, and J Xu. Natural movement: or, configuration and attraction in urban pedestrian movement. *Environment and Planning B: Planning and Design*, 20(1):29–66, 1993.

[18] Rebecca Hurlebaus, Kai Basten, Hanspeter A. Mallot, and Jan M. Wiener. Route learning strategies in a virtual cluttered environment. In Christian Freksa, Nora S. Newcombe, Peter Grdenfors, and Stefan Wlfl, editors, *Spatial Cognition VI. Learning, Reasoning, and Talking about Space*, volume 5248, pages 104–120. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[19] Google Inc. Google maps API family. http://code.google.com/apis/maps/.

[20] Petra Jansen-Osmann and Petra Fuchs. Wayfinding behavior and spatial knowledge of adults and children in a virtual environment. *Experimental Psychology (formerly "Zeitschrift fr Experimentelle Psychologie")*, 53(3):171–181, January 2006.

[21] Gabriele Janzen and Miranda van Turennout. Selective neural representation of objects relevant for navigation. *Nature Neuroscience*, 7(6):673–677, May 2004.

[22] Panayotis Kikiras, Vassileios Tsetsos, and Stathes Hadjiefthymiades. *Ontology-based user modeling for pedestrian navigation systems*, pages 1–6. 2006.

[23] Benjamin Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2(2):129–153, 1978.

[24] Benjamin Kuipers, Dan G. Tecuci, and Brian J. Stankiewicz. The skeleton in the cognitive map: A computational and empirical exploration. *Environment & Behavior*, 35(1):81–106, January 2003.

[25] Sharon A Livingstone and Ronald W Skelton. Virtual environment navigation tasks and the assessment of cognitive deficits in individuals with brain injury. *Behavioural Brain Research*, 185(1):21–31, December 2007. PMID: 17727970.

[26] Kevin Lynch. *The image of the city*. M.I.T. Press, Cambridge Mass. [u.a.], 1960.

[27] T. Meilinger and M. Knauff. Ask for directions or use a map: A field experiment on spatial orientation and wayfinding in an urban environment. *Journal of Spatial Science*, 53(2):13–23, 2008.

[28] Tobias Meilinger, Gerald Franz, and Heinrich H Blthoff. From isovists via mental representations to behaviour: first steps toward closing the causal chain. *Environment and Planning B: Planning and Design*, 0(0), 2011.

[29] Pierre-Emmanuel Michon and Michel Denis. When and why are visual landmarks used in giving directions? In Daniel R. Montello, editor, *Spatial Information Theory*, volume 2205, pages 292–305. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001.

[30] R. Passini, H. Pigot, C. Rainville, and M.-H. Tetreault. Wayfinding in a nursing home for advanced dementia of the alzheimer's type. *Environment and Behavior*, 32(5):684–710, September 2000.

[31] Roy A Ruddle and Simon Lessels. Three levels of metric for evaluating wayfinding. *Presence: Teleoperators and Virtual Environments*, 15(6):637–654, December 2006.

[32] Roy A Ruddle, Ekaterina Volkova, Betty Mohler, and Heinrich H Blthoff. The effect of landmark and body-based sensory information on route knowledge. *Memory & Cognition*, 39(4):686–699, May 2011. PMID: 21264583.

[33] A W Siegel and S H White. The development of spatial representations of large-scale environments. *Advances in Child Development and Behavior*, 10:9–55, 1975. PMID: 1101663.

[34] Molly E. Sorrows and Stephen C. Hirtle. The nature of landmarks for real and electronic spaces. In *COSIT 99. VOLUME 1661*, pages 37–50. Springer, 1999.

[35] Brian J Stankiewicz and Amy A Kalia. Acquisition of structural versus object landmark knowledge. *Journal of Experimental Psychology. Human Perception and Performance*, 33(2):378–390, April 2007. PMID: 17469974.

[36] Perry W Thorndyke and Barbara Hayes-Roth. Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14(4):560–589, October 1982.

[37] Barbara Tversky. Some ways that maps and diagrams communicate. In *Spatial Cognition II, Integrating Abstract Theories, Empirical Studies, Formal Methods, and Practical Applications*, pages 72–79, London, UK, 2000. Springer-Verlag.

[38] Barbara Tversky and Paul U. Lee. Pictorial and verbal tools for conveying routes. pages 51–64. Springer, 1999.

[39] S. Werner, B. Krieg-brckner, and Herrmann. Modelling navigational knowledge by route graphs. In *In: In: Ch. Freksa et al. (Eds): Spatial Cognition II 1849*, page 295316. Springer, 2000.

[40] Janis Williamson and Corinne Barrow. Errors in everyday routefinding: A classification of types and possible causes. *Applied Cognitive Psychology*, 8(5):513–524, October 1994.

[41] F. B. Zhan and C. E. Noon. Shortest path algorithms: An evaluation using real road networks. *Transportation Science*, 32(1):65–73, February 1998.

[42] Don H. Zimmerman and D. Lawrence Wieder. The diary. *Journal of Contemporary Ethnography*, 5(4):479–498, 1977.

Glossary

cognitive map
A mental map of the environment formed by humans to enable determination of location and creation of routes.. 1.0

decay
Loss of information within the user memory over time. Indicates the deterioration or decline in stored data.. 2.4

deformed grid
Grid with a number of paths and nodes removed to produce a more complex, but still fully accessible, representation of the environment.. 2.2

destination point
Location which is the aim for wayfinding.. 2.2, *see also* end point

directions
Series of humanly readable instructions which can be used to traverse a route.. 2.2

end point
Location which is the aim for wayfinding.. 2.3, *see also* destination point

environment representation
Model of the environment to be traversed.. 3.1

familiarity
How regularly a route component has been visited and has been stored in memory. Gives an indication of how well known the component is to the user.. 1.0

fewest turns
Route criterion which minimises the number of turns in the route. *see also* Minimum Turns

forgetting
Loss of detail during the process of transferring route information through directions.. 2.4

grid

Basic, symmetrical, representation of the environment in which nodes are fully connected by paths to each of their neighbours. Graph representation of the environment.. 2.2

inaccuracy

Errors introduced by loss of concentration or misconception of location.. 4.3

landmark

Visual reference point used to indicate that the user is travelling in the right direction or indicating where a change of direction is required.. 1.0

landmark strength

Combination of the prominence, singularity, prototypicality and content of a landmark, along with its distance from the current location.. 4.2

link

Section of a path between two nodes.. 2.2

Longest Leg First

Similar to Minimum Turns, but the route should have the longest feasible distance before the first turn.. 2.3

metric

Some specific measure of the performance of a system or algorithm.. 2.5, *see also* route metric

node

Decision point on the route. Point at which the user can potentially choose more than one path. Junction or intersection between multiple routes.. 2.2

path

A trail, footpath, road or similar feature along which a pedestrian may travel.. 2.2

pointer

Variable which contains the address of an object in C.. 3.2

recognition factor

Single value formed from the familiarity of the current component, any attached components and landmark strengths where appropriate.. 4.3

route

Sequence of route components, which when travelled through in order, will indicate how to traverse from the start point to the end point.. 1.0, 3.2

route component
: Node, Link or Landmark which can be strung together in a sequence to create a route.. 2.4, 3.2

route metric
: Some specific measure of a route attribute which allows for comparison between routes. Examples of this are route length and route complexity.. 2.5, *see also* metric

Shortest Path
: Route between two nodes in a grid which, of all possible routes, has the least cost in terms of distance.. 1.0

start point
: Location at which the user begins wayfinding.. 2.2

user memory
: Storage for route components and routes known to the user or the virtual user.. 4.0

user simulation
: Algorithm to approximate the behaviour of a human subject when traversing a route.. 2.1

wayfinding
: Movement between the start point and destination point, where the latter is not visible from the former.. 1.0

Appendix A

Personal Reflection

This project has posed one of the biggest challenges I have ever had to face, but it has also brought me a huge amount of pleasure. By managing to remain enthusiastic about the topic throughout the process, I have achieved many of the things I intended and kept motivated even through the difficult times. I think that the most important thing I have learnt during this exercise is that it is as important to know when to stop as it is to know where to start. Researching previous work, designing system elements or even increasing the complexity of the functionality can be essential, but knowing when to switch between tasks is also crucial. If one of these exercises is allowed to overrun or become extremely intricate without consideration of how it will be evaluated, it can affect all aspects of the project still to be completed, reducing available time and increasing the effort required.

In a similar way, it is easy to drift away from the goal at a tangent, or produce huge amounts of code that is irrelevant. Getting regular supervisor input can help to avoid this, along with a well thought out methodology, and referring to the background research and project aims routinely during implementation. Even if this does happen, it is important not to panic and to realise that if the code is never going to fit the task, it is better to throw it away than spend many hours trying to force it to work.

Wherever possible, break the implementation down into smaller manageable tasks, keeping each one as simple as possible. By doing this time can be allocated easier, and any required modifications can be completed with minimum fuss and disruption. Test frequently, between each of these stages if possible, as locating and fixing errors is simpler if they are caught early before the incorrect values or concepts are allowed to travel through the entire system. Mistakes not found when they occur can produce a ‘house-of-cards’ effect where later code relies on these errors, and correcting them involves replacing whole sections of code.

Allow plenty of time for evaluation, especially when building a system from scratch. If there are no previous results or datasets to compare the performance of your algorithms to, then selecting suitable metrics and producing enough information to establish quality can be difficult and time consuming. Also, even the best written code can behave erratically when being tested on a model other than the one it was developed using, so it may be necessary to make many changes to further investigate or justify these observations.

Finally, the write-up will always take longer than anticipated. Start this well ahead of time and get plenty of feedback from your supervisor, they have written or read many papers and theses so their advice can be invaluable.

Appendix B

A Brief Record of Materials Used in the Solution

No external code or datasets were used in the production of this system. All work is original to the author, and no materials, drafts or notes were provided in advance of its start or completion.

Appendix C

Data Collection

This appendix contains the information pertaining to the collection of suitable data from a number of users. Firstly a map of the campus is shown, followed by the pedestrian route survey employed, and finally a table containing the data resulting from this process.

C.1 Campus Map

Figure C.1: Detailed campus map (taken from http://www.meetinleeds.co.uk/pdfs/Colour_Campus_Map.pdf).

C.2 Pedestrian Route Survey

Aim

The aim of this survey is to collect data for a project based on pedestrian routes around the University campus. Any information provided should be on a voluntary basis, and all surveys will be completely anonymous.

Instructions

The following three pages contain tables and a map which should be filled in in the following way:
Start point – this should be the building and room or area from which the journey originated.
End point – this should be the building and room or area at which the journey ended.
Route – please give a brief description of the route taken between the two points.
Please indicate the route on the map provided on the final page if possible.

Examples

Start point – MSc study, EC Stoner.
End point – Student Union shop.
Description – Out through the door by the study, up the steps past the Edward Boyle Library, under the walkway outside Social Sciences. Left at the end and through the gateway. Down the steps at the front of the union and in through the door by Santander.

Notes

Journeys undertaken entirely within the same building do not need to be entered. Details of movements once inside a building are not required. Please provide details of any journeys undertaken on the University of Leeds campus for the whole of one day.

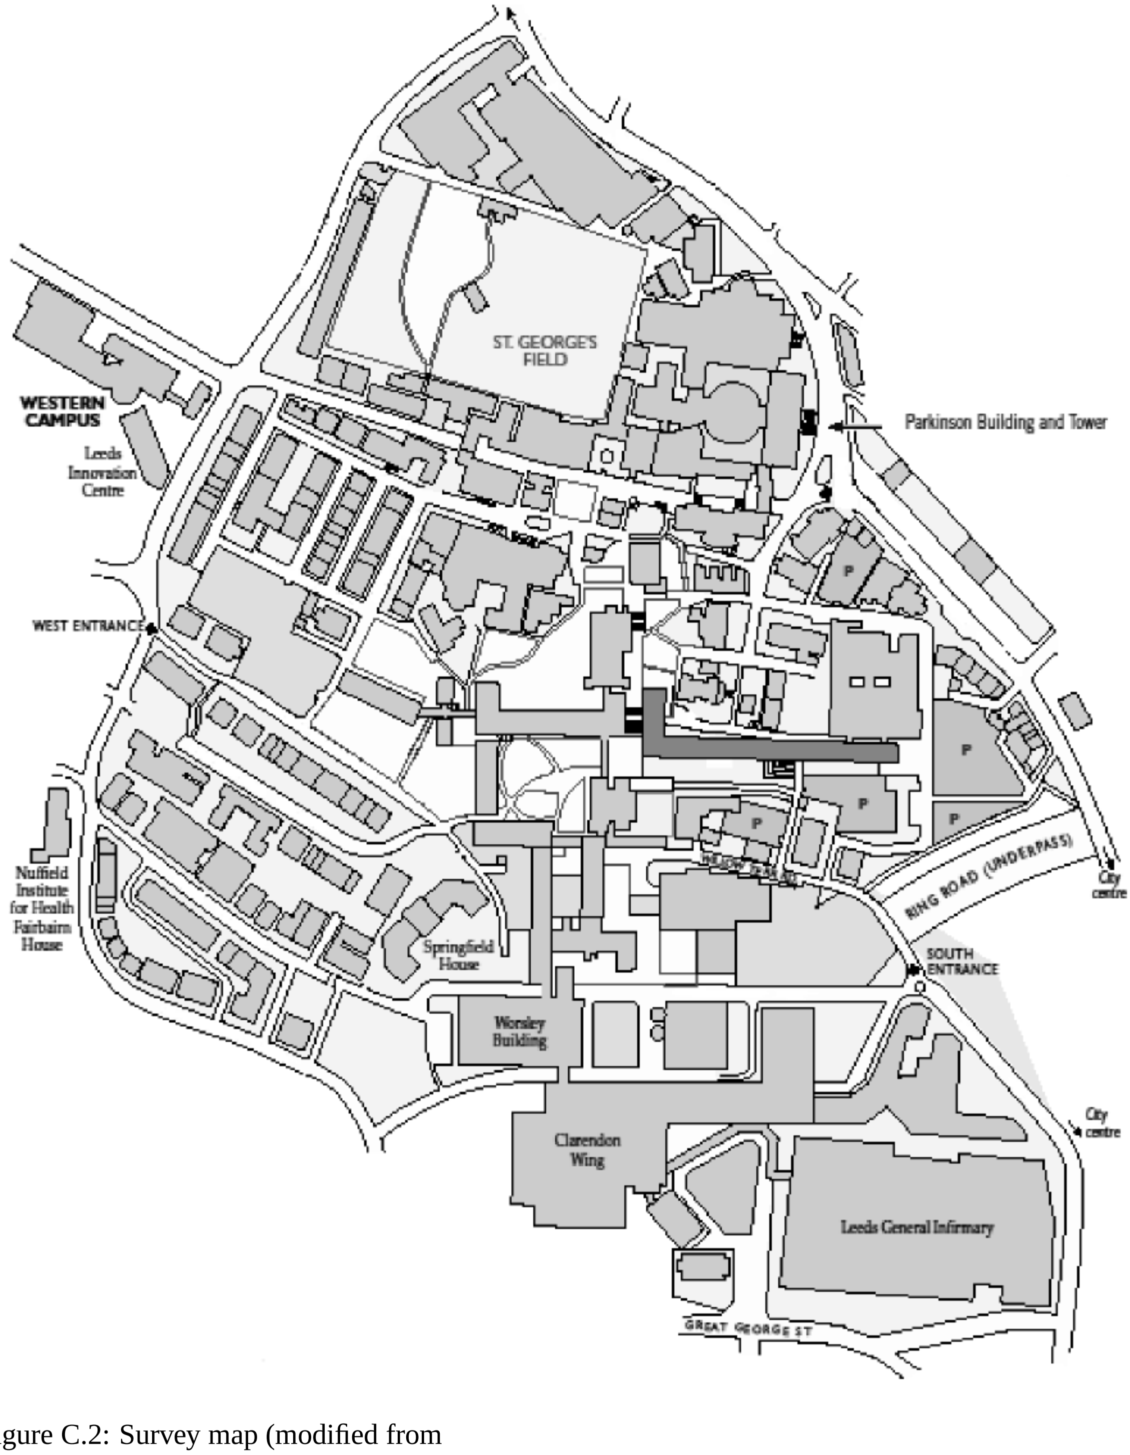

Figure C.2: Survey map (modified from
http://www.engineering.leeds.ac.uk/faculty/contact/documents/map.pdf).

		Start Point	Route	End Point
1	Bldg Room			
2	Bldg Room			
3	Bldg Room			
4	Bldg Room			
5	Bldg Room			
6	Bldg Room			
7	Bldg Room			
8	Bldg Room			
9	Bldg Room			
10	Bldg Room			

Table C.1: Survey diary page.

C.3 Collected Data

Participant 1	EC Stoner MSc Study	Out through the door by the study, up the steps past Edward Boyle Library, under the walkway outside Social Sciences, right at the end and through the automatic door opposite the toilet.	Parkinson Building B11
	Parkinson B11	Past the toilets, through the automatic door, past the Michael Sadler Building and through the gateway. Through the entrance to the union and down the steps.	Student Union Shop
	Student Union Lounge	Out through the door, left from the union and left again along Cromer Terrace. Right and through the entrance	Careers Centre
	Careers Centre	Out through the door, left along Cromer Terrace, right at the end, past the union and through the gateway. Through the walkway, down the steps past Edward Boyle and through the door into the study.	MSc Study
	EC Stoner MSc Study	Out through the door , up the steps near Edward Boyle, under the walkway outside social sciences, left at the end, past the union, school of psychology and business school and then left and through the entrance.	New Law Building
Participant 2	Charles Mo	Out gates, around Whetton, down stairs and through courtyard	EC Stoner
	EC Stoner	Exit from red route, tight left then up around the side of building	Union
	Union	Out door straight and then follow alley through Chem building and Workers court	Elec Eng
	Elec Eng	Out door, through Chem block alley back through Worker's court	Union
	Union	Out back door, across green to gates	Charles Mo
Participant 3	MSc Study	Out the door by the study, turn left, down the stairs (near the velo campus), walk straight ahead, turn right, take the stairs to the Edge	the Edge
	the Edge	Out the stairs, straight on ahead, under the EC Stoner, turn left, then right and straight ahead to end point.	The Parkinson Building
	Parkinson Building	Out the building by the main gates and straight to Ziff Building	Marjorie & Arnold Ziff Accommodation Office
	Ziff Building	Out the doors, right and take the road right next to the Ziff building, take the stairs in front of Edward Boyle Lib, take a left and up the stairs to the Union.	Union
	Union	Out the doors, take a right, straight on to the red corridor, at the end of of the red corridor turn right to the cafe	EC Stoner Dol Ce Vita Cafe
	EC Stoner Dol Ce Vita Cafe	Straight out the Edward Boyle, take a left and down the stairs (by School of Mathematics), right at the entrance of Roger Stevens, walk up 1 level to the ISS.	Roger Stevens Irene Manton ISS cluster
	Roger Stevens Irene Manton ISS cluster	Out by the main doors of Roger Stevens, walk straight on ro reach the union, then take a left at the end of the road and walk straight and a right	Careers Centre

Table C.2: Collected route data: Participants $1 - 3$.

Participant 4	Entrance	Enter through south entrance, follow road to Edge entrance. Take right turn, cut across the car park. Enter EC Stoner staircase 1.	EC Stoner MSc Study
	EC Stoner MSc Study	Out door by study, up stairs by library. Under cover by Social Sciences, turn right. Down slope under walkway, in through door.	Parkinson Building B11
	Parkinson Building B11	Out of door, turn right, walk down past Ziff Building follow path straight on, in through SOC staircase 1	EC Stoner MSc Study
	EC Stoner MSc Study	Out door by study, up steps and round past the front of the Union. Follow the path and then the road. Turn left into Clarendon Place, down the road and in.	Student Counselling
	Student Counselling	Out and up Clarendon Place. Turn right and straight past the front of the Union. Turn right and go down the steps and in through the door.	EC Stoner MSc Study
Participant 5	EC Stoner MSc Study	Out of the door nearest the study up the big steps past equality service. Turn right and follow to brown door in Parkinson Building	Parkinson Building B11
	Parkinson Building B11	Out of door on south side of Parkinson Building down steps. Follow road down past front of music dept. Across to EC Stoner entrance 1 north.	EC Stoner MSc Study
	EC Stoner MSc Study	Out of door nearest MSc Study. Diagonally across to furthest door in side of RS Building	Roger Stevens LT8
	Roger Stevens LT8	Along red route and follow to come out near equality service, carry straight on, turn left and go past Union buildings, turn right then left, turn right entrance to George's field. Cross George's field on left hand path. Cross road and right to rotating doors	Engineering Houldsworth LTB
	Engineering Houldsworth LTB	Out and left. Cross George's field. Past Geography to main part and through arch with crest and past old buildings. Turn right, past grass then left and equality service roughly ahead.	Equality Service
	Equality Service	Out of door turn left down steps left to study door	EC Stoner MSc Study
	EC Stoner MSc Study	Up big steps, left just before end into side door of old bar	Old Bar Union Building
	Old Bar Union Building	Out of side door of old bar across where there's lots of steps and along to lower door of Parkinson Building	Parkinson Building Brotherton Library
	Parkinson Building Brotherton Library	Out of Parkinson Building onto steps. Turn left follow road which goes past music. Cut across to EC Stoner and in entrance 1 north	EC Stoner MSc Study

Table C.3: Collected route data: Participants 4 and 5.

Appendix D

Project Plans

This appendix contains the Gantt charts showing the original and final (ammended) project schedules.

D.1 Original Project Plan

Task Name	Start	Finish
Background Reading	25/02/11	17/06/11
Data Collection	28/03/11	08/04/11
Interim Report Creation	01/06/11	16/06/11
Interim Report	17/06/11	17/06/11
Chapter Creation	11/07/11	21/07/11
Chapter Submission	22/07/11	22/07/11
Report Writing	29/05/11	31/08/11
Basic User Simulator	31/05/11	13/06/11
Extended User Simulator	17/06/11	21/06/11
User Simulator Completion	21/06/11	22/06/11
Route Creation Algorithm Design	22/06/11	29/06/11
Basic Route Creation Algorithm Implementation	01/07/11	15/07/11
Testing	16/07/11	23/07/11
Basic Route Creation Algorithm Completion	23/07/11	24/07/11
Advanced Route Creation Algorithm Implementation	24/07/11	29/07/11
Progress Meeting	29/07/11	30/07/11
Testing	01/08/11	07/08/11
Advanced Route Creation Algorithm Completion	07/08/11	08/08/11
Evaluation	08/08/11	14/08/11
Final Report Draft	19/08/11	31/08/11
Report Submission	01/09/11	01/09/11

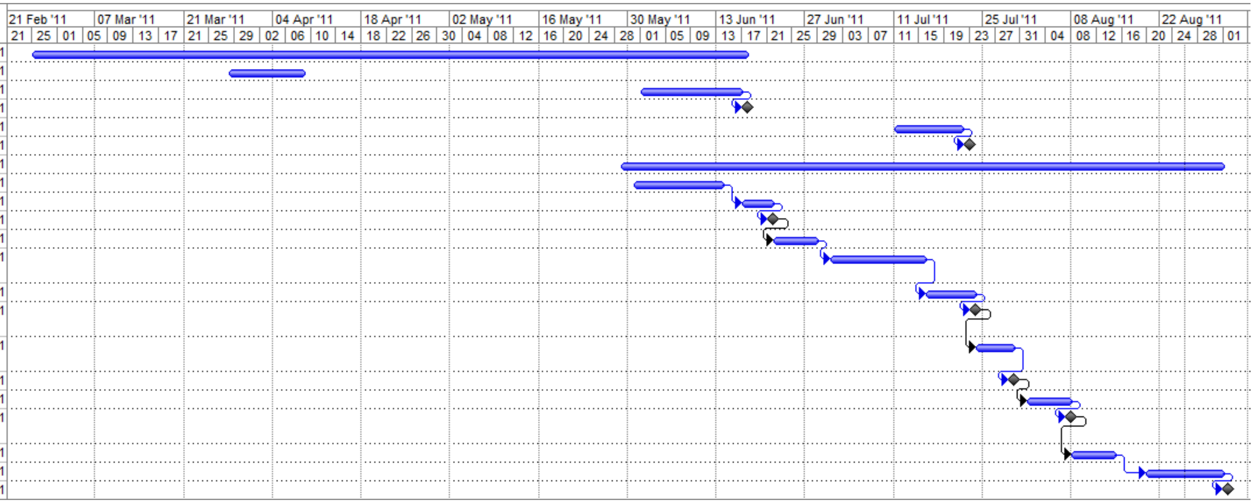

Task Name	**Start**
Interim Report	17/06/11
User Simulator Completion	21/06/11
Chapter Submission	22/07/11
Basic Route Creation Algorithm Completion	23/07/11
Progress Meeting	29/07/11
Advanced Route Creation Algorithm Completion	07/08/11
Report Submission	01/09/11

D.2 Final Project Plan

Task Name	Start	Finish
Background Reading	25/02/11	17/06/11
Data Collection	28/03/11	08/04/11
Interim Report Creation	01/06/11	16/06/11
Interim Report	17/06/11	17/06/11
Chapter Creation	11/07/11	21/07/11
Chapter Submission	22/07/11	22/07/11
Report Writing	29/05/11	31/08/11
Basic User Simulator Including Basic Environment	31/05/11	13/06/11
Testing	13/06/11	14/06/11
Extended User Simulator Including Extended Environment	17/06/11	23/06/11
Testing	23/06/11	24/06/11
User Simulator Completion	24/06/11	24/06/11
Overall System Design	24/06/11	27/06/11
Route Creation Algorithm Design and Implementation Including User Memory	27/06/11	05/07/11
Testing	05/07/11	10/07/11
Route Creation Algorithm Completion	11/07/11	11/07/11
Creation of the Real Environment and Extensions	12/07/11	20/07/11
Extended Algorithm Completion	21/07/11	21/07/11
Evaluation Phases 1 - 3	22/07/11	28/07/11
Progress Meeting	29/07/11	29/07/11
Evaluation Phases 4 & 5	30/07/11	11/08/11
Evaluation Completion	12/08/11	12/08/11
Final Report Write-up	12/08/11	31/08/11
Report Submission	01/09/11	01/09/11

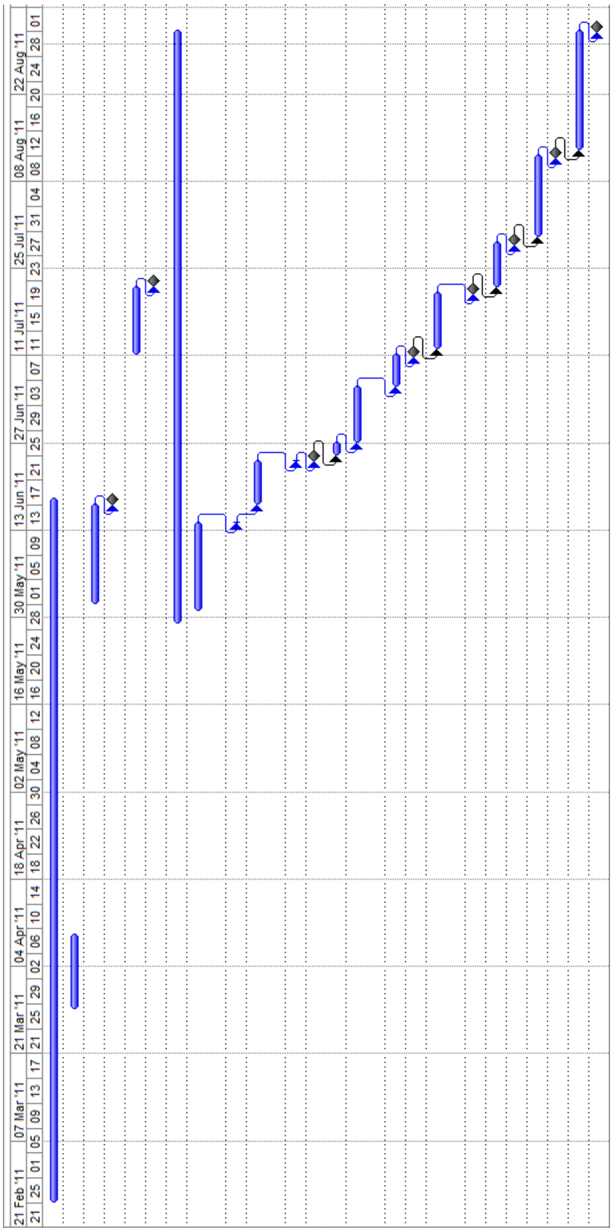

Task Name	**Start**
Interim Report	17/06/11
User Simulator Completion	24/06/11
Route Creation Algorithm Completion	11/07/11
Extended Algorithm Completion	21/07/11
Chapter Submission	22/07/11
Progress Meeting	29/07/11
Evaluation Completion	12/08/11
Report Submission	01/09/11

Appendix E

Interim Report