Common ground breakdown during collaborative virtual environment navigation with wall-sized and desktop displays

Jeremy Andrew Swann

Submitted in accordance with the requirements for the degree of
Doctor of Philosophy

The University of Leeds
School of Computing

September 2015

Acknowledgements

I would like to take this opportunity to thank a number of people without whom this research wouldn’t have been possible.

Firstly, I would like to thank my supervisor Dr Roy Ruddle for the years of excellent support both academically and personally. Also, the feedback and advice from Dr David Duke and Dr Vania Dimitrova when reviewing my PhD has been gratefully received.

To my fellow colleagues in the Computational Science and Engineering (CSE) group, you’ve made coming in every day something to look forward to. Special mention must go to Trien Van Do, John Hodrien and Rhys Thomas for helping me throughout. I must also thank the many others I know in the department, Leeds and home for understanding and supporting me throughout my PhD.

Lastly, but by no means least, my thanks go to my family for their ever present support and interest in everything I do.

This research was funded by a Doctoral Training Grant Studentship from the School of Computing

Abstract

During collaborative data analysis participants are reliant on a shared common ground to be able to understand others' locations and actions, breakdowns in which require remedial action before other activities can continue. This thesis investigates the frequency and severity of common ground breakdowns that occur during collaborative navigation, when a wall-sized display user (master) directs a desktop user (slave) within a Collaborative Virtual Environment (CVE).

A series of experiments were conducted to investigate this form of collaborative navigation and evaluate the developed solutions. These experiments required participants using a desktop display to view a wall-sized display user's navigation to targets within a 3D landscape, before attempting to re-visit these targets themselves.

Experiments 1 and 2 were conducted to establish the typical frequency and severity of common ground breakdowns when the desktop user is attempting to find singular and multiple targets. Participants exhibited non-trivial levels of common ground breakdown over different types of movement and input device used by the wall-sized display master. Although they frequently had sufficient common ground to reach the approximate area of a target, this was insufficient for them to be able to complete the task.

The remainder of the research investigated two distinct classes of solution to these breakdown: additional views and path visualisations. For additional views, a large field of view (FOV) context view and local overview map were provided. Experiment 3 evaluated their effect and demonstrated that neither of which reduced the level of common ground breakdown exhibited by desktop users. Behavioural changes of participants using the context view still led to similar task failures.

For path visualisation, two representations of paths between targets, string and heatmap were provided. Experiment 4 provides evidence that both representations significantly increased participant's success rate. Behavioural data showed that par-

ticipants visited areas unrelated to their search less frequently and remained closer to the desired path.

In conclusion, this research has four major contributions. First, a classification of the types, severity and frequency of common ground breakdowns that occur between desktop and wall-sized display users conducting master-slave navigation within a CVE. Second, evidence is provided that additional views are not beneficial in reducing the level of breakdown for the desktop users. Third, path visualisations are shown to be effective in reducing the level of breakdown experienced by the desktop users, and allowing more effective navigational behaviour. Finally, the successful application of heatmaps in aiding navigation when previous applications have been limited to analytical use.

Abbreviations

CAVE	Cave automatic virtual environment
CSCW	Computer-supported cooperative work
CSV	Comma separated values
CVE	Collaborative virtual environment
DEM	Digital Elivation Model
FOV	Field of view
HMD	Head mounted display
ICC	Institut Cartogrfic de Catalunya
VE	Virtual environment
VR	Virtual reality

Contents

List of Figures

Chapter 1

Introduction

Modern computing technology is being used in a wide variety of contexts from portable mobiles and tablets, through ”conventional” laptops and desktops, to systems that output to larger displays which occupy entire walls or rooms. The large display size and resolution of tiled, wall-sized displays combined with a user’s ability to physically navigate the data visible upon it allows navigation with greater speed and accuracy than desktop displays [9]. This has been utilised in areas such as digital pathology, collaborative sense-making and life sciences to enable users to investigate large, multi scale datasets effectively [55, 122, 124].

However, the complex nature of the datasets that wall-sized displays are most suitable for often needs to be analysed by multiple specialists, who will likely be distributed across different physical locations. The limited availability of these forms of display (due to cost, space or portability) means that some of the individuals within this collaborative work will be utilising a more modest desktop display. These location and display differences add to other differences between users that contribute to disruptive breakdowns to the common ground between users that is required to conduct collaborative tasks [28].

This thesis investigates the occurrence and severity of common ground breakdown between desktop display and wall-sized display users through a distributed geological scenario in a Collaborative Virtual Environment (CVE) (see figure 1.1). This includes the adaptation of navigational aids, previously effective in desktop-only collaboration to reduce these breakdowns to a sufficient level to allow for beneficial collaboration between users.

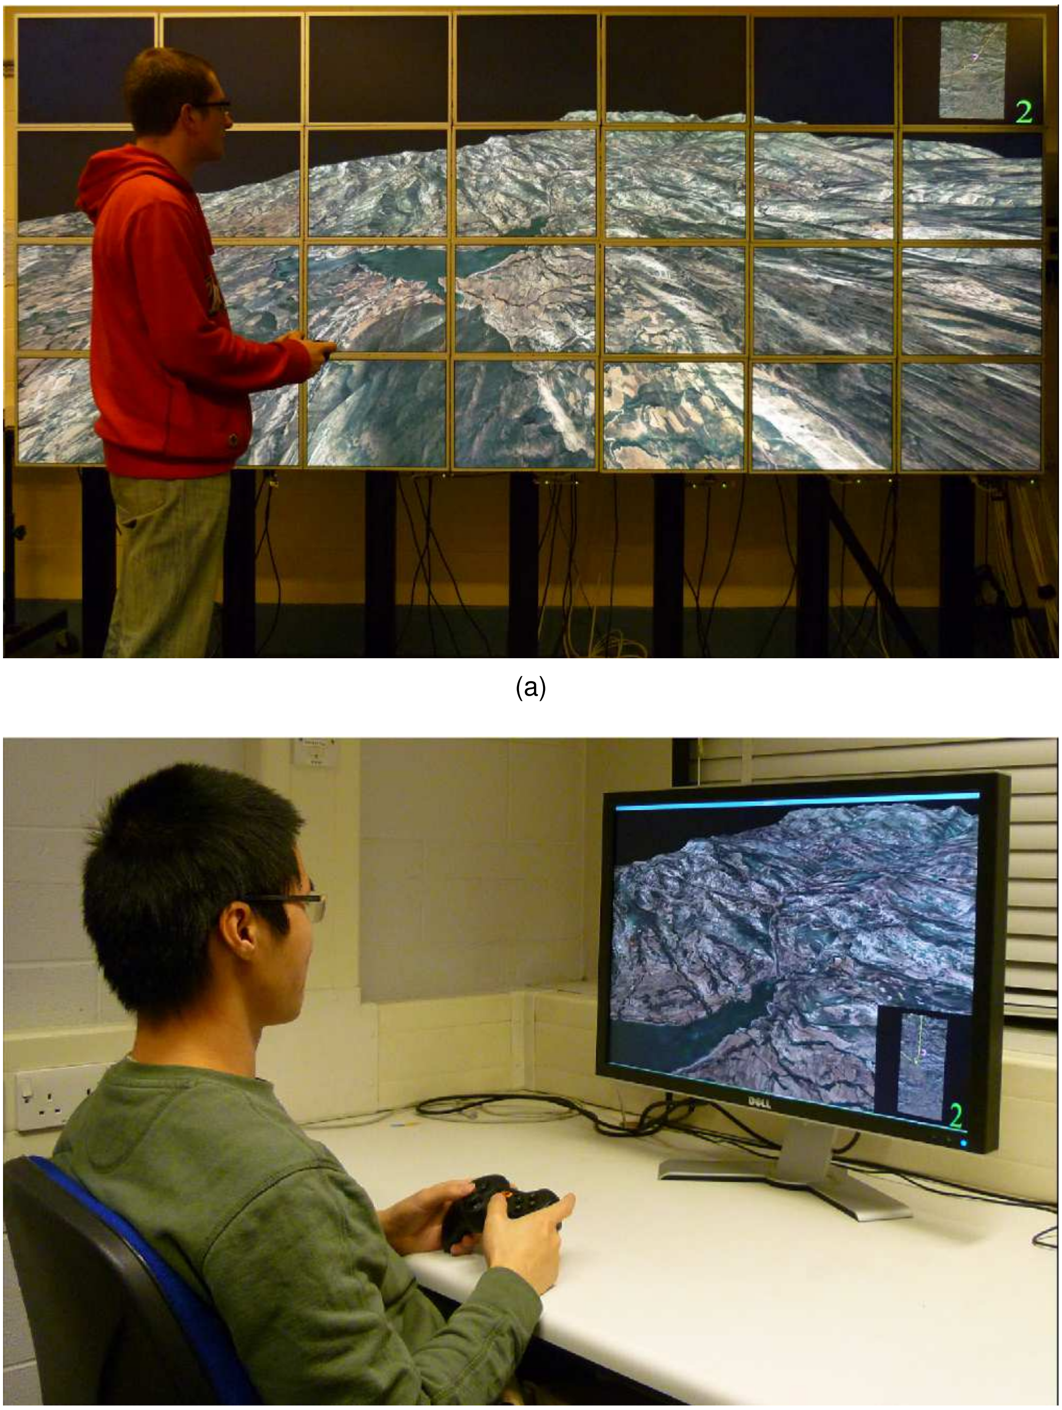

(a)

(b)

Figure 1.1: Collaborators using a CVE to navigate a geological dataset: (a) Master user with a 54 million pixel, wall-sized display and (b) Slave user with a desktop display.

1.1 Thesis outline

This chapter introduces the primary concepts and contributions that will be explored within the rest of the thesis. Chapter 2 details the topics and previous work that informs this research. It initially discusses the importance of common ground in conversational analytics and wider collaborative scenarios. Navigational strategies and evaluation methods are introduced before presenting the difficulties users have in navigating with CVEs and the aids used to alleviate these difficulties in desktop CVEs. Following from this, different display arrangements such as tiled, wall-sized displays, cave automatic virtual environments (CAVEs) and tablets are introduced alongside the issues when collaborating across these different displays. A suitable geographic scenario is introduced as well as the types of common ground breakdowns that will occur. Lastly, evaluation methods that will be used throughout the research are discussed.

Chapter 3 details the CVE system which was developed to allow the research described in later chapters to be conducted. This CVE allows the navigational movements of users to be recorded, replayed and analysed. The format of the dataset used to generate the environment, as well as the input and output devices utilised, are described before presenting examples of the logs that are generated from the system.

A series of experiments were conducted in chapters 4, 5 and 6 to investigate common ground breakdowns within collaborative navigation between a wall-sized master and desktop slave and evaluate the solutions developed. These experiments required participants using a desktop display to view a wall-sized display user's navigation to targets within a 3D landscape, before attempting to re-visit these targets themselves.

Chapter 4 presents the first two experiments which were conducted to establish the typical types, frequency and severity of common ground breakdowns when the desktop user is attempting to find singular and multiple targets. Both experiments present the method used before presenting results on successful and failed trials alongside typical behaviour exhibited for each. The findings of each experiment are discussed before being concluded at the end of the chapter.

The remainder of the research investigated two distinct classes of solution to these breakdowns: additional views and path visualisations. Chapter 5 presents two types of additional views, a large field of view (FOV) context view and a local overview map. Experiment 3 evaluated their effect and demonstrates that neither reduced the level of common ground breakdown exhibited by desktop users. Behavioural changes of participants using the context view still led to similar task failures.

Chapter 6 investigates path visualisations, of which two representations were provided, string and heatmap. Experiment 4 provided evidence that both representations significantly increased participants' success rate. Behavioural data showed that participants visited areas unrelated to their search less frequently and remained closer to the desired path.

Chapter 7 concludes the research by summarising and assessing the main contributions before discussing potential future work that can utilise the presented findings.

1.2 Contributions

The research in this thesis has resulted in the following contributions to the research area:

1. A classification of types, severity and frequency of common ground breakdowns that occur between desktop and wall-sized display users conducting master-slave navigation within a CVE. This classification was established through a scenario where participants were conducting closely coupled tasks consisting of revisiting targets previously seen through distributed and asynchronous collaborative navigation. Completing these tasks required freeform navigation (x/y/z translational and heading/ pitch movement) of a large environment (42.5 x 74.6 kilometre) which provided high optic flow but few suitable navigational landmarks.

2. Evidence that additional views, a type of aid previously beneficial in desktop navigation, do not reduce the frequency or severity of common ground break-

downs within the current scenario. Participants' navigational behaviour indicated that the addition of a wide FOV context view prompted participants to adopt a route-based navigation strategy more frequently, even though this ultimately led to similar navigational failures occurring.

3. Evidence that another type of aid, heatmap and string-based path visualisations, reduce the frequency of breakdowns experienced by desktop users. Both heatmap and string visualisations resulted in participants visiting areas unrelated to their search less frequently and remaining closer to the desired path.
4. The successful application of heatmaps in aiding navigation when previous use has been predominately limited to analytical use. Use of this aid resulted in similar success rates and quicker search performance when compared to the more frequently used string-based path visualisation.

Chapter 2

Background

This chapter details the topics and previous work that informs the research within this thesis. It begins by introducing common ground as part of conversational analytics and wider collaborative scenarios. Ways to classify these wider collaborative scenarios are considered as part of computer-supported cooperative work (CSCW). The navigational strategies used by humans are detailed as well as ways to evaluate this navigation to understand the performance, behaviour and rationale behind one's actions. The difficulties of Collaborative Virtual Environment (CVE) navigation are introduced before presenting typical navigational aids that are used to alleviate these difficulties in desktop CVEs. Following from this, different display arrangements such as tiled, wall-sized displays, cave automatic virtual environments (CAVEs) and tablets are discussed and the issues when collaborating across these different displays presented.

Next, a geographic scenario is introduced alongside the design choices for its use within a wall-sized to desktop display CVE. After this, the sources of common ground breakdown within this scenario are presented alongside examples of functionality that have previously been used to address similar breakdowns. Lastly, the evaluation methods that will be used are discussed.

2.1 Common Ground

Common ground is the mutual knowledge, beliefs and assumptions between two (or more) people that is required to conductive collaborative activities. Conversational analysts such as Clark and Brennen [28] state that it is required to allow people to coordinate the process of a collaboration (e.g., what language to speak, cultural norms to follow) and needs to be updated on a moment by moment basis (e.g., adjusting the appropriate level of one's voice depending on a hearing disability). Only when this is in place can they begin to coordinate content, the reason why an activity is conducted in the first place.

2.1.1 Grounding

Updating common ground requires a process called *grounding* and involves seeking confirmation that assumed common ground is correct and, when proven otherwise, the repair of incorrect common ground [28]. The loss of a mutual understanding of what is being spoken about (content common ground), can be described using the term breakdown [6]. This phrase can describe circumstances which causes a change from ”routine practice to problem solving”, resulting in coping behaviour and other workaround actions [66].

Clark and Brennen stated that conversationalists seek to repair faults before they propagate further [28], suggesting that breakdowns are repaired immediately before other activities can continue. However, people seek to utilise the minimum amount of effort possible to complete many tasks [29, 58], so the purpose and context of the communication impacts the level of breakdown at which participants dictate repair is required and the form that these actions take [28]. A casual conversation will result in common ground only being repaired when the participants believe there are large differences in their common grounding, whereas the discussion of a legally binding document will require a very high level of common ground, with almost constant grounding and repair processes occurring to ensure there are no misunderstandings.

The communication medium that is used to conduct a collaboration also affects the type of actions that can be performed by participants. For example, using a telephone allows instantaneous reception and transition of speech, which will be used differently to the delayed, text based delivery of email messages. Clark and Brennen [28] classified these differences according to 8 constraints which, in turn, affect the costs associated with starting, producing, receiving and understanding different actions.

The levels of these costs can be used to help understand differences in participants' actions due to the medium being used, such as pre-empting potential uncertainties in an email response which could be corrected as they occur when using the telephone. They can also be used to consider how changing elements of the communication medium may affect the collaboration. For example, switching from email to instant messaging may result in a change to more dynamic correction of misunderstandings.

2.1.2 Breakdowns

Although these concepts are derived from conversation analysis, they can be applied to other collaborative activities where participants must have sufficient common ground in the process and content of that activity. A simple example is dancing requiring knowledge of the steps (content) and tempo (process) of a particular dance.

A more relevant example of collaborative work is presented by Convertino et al [31, 32, 33], who investigated common ground within an emergency management planning scenario. This work emphasised that both process and content common ground are important for a successful collaboration as the establishment of process common ground allowed content common ground to be shared, understood and expanded more effectively.

Different elements of such a scenario affect members' ability to establish, maintain and repair different forms of common ground. Here, users have to establish common ground with a group that they have met for the first time under significant time pressures, in a distributed environment that removes many of the shared cues found in a collocated environment. This task is made more difficult with members coming from differing areas of expertise that have their own roles, skills and language that can cause misunderstandings between different members.

When areas of potential breakdown have been established, an appropriate way of measuring their effect within users' activities must be determined. In this scenario, the complexity and variety of activities means the authors measured breakdowns through analysis of the number and type of conversation turns (changes in the current topic).

2.2 CSCW

As the types of common ground and associated breakdown vary greatly depending on the specific scenario, it must be possible to accurately and comprehensively define this before it can be addressed. Computer Supported Cooperative Work (CSCW) provides classifications that allow these collaborative scenarios to be understood and the scope of specific activities accurately defined. When this has been defined, previous examples of similar scenarios can be considered to provide supportive systems to aid users.

Collaborative activities occur constantly, from individuals executing their own activities with only a passing consideration of what effect their actions have on the rest of a group or organisation, to multiple people simultaneously working together within a single activity. These activities can be classified by the time/space matrix, group membership/ processes, combined classifications or system functionality.

2.2.1 Time/space matrix

One of the most basic aspects of a collaborative scenario is the space and time in which it occurs [74]. Space may be collocated, where the users are in close physical proximity, or distributed, where they are geographically distant. Time may be synchronous, where the collaborators are interacting simultaneously, or asynchronous where this interaction occurs at different periods in time. Collaboration in the same physical location provides collaborators with a large amount of context from the shared location and timing of their interaction [84]. When this collaboration occurs in a distributed setting a number of the previously common elements of collaborators' context are lost and requires the group to explicitly acquire this through different group processes or additional support through the systems they are using (e.g., summarising important elements of a user's view of the data to other collaborators [27, 30]).

2.2.2 Group membership and processes

The length, formality and size of group membership can also vary greatly, but Andriessen has simplified them to three settings of collections, knowledge sharing communities and teams [2]. Collections are loosely connected groups of people who collaborate for a short, undefined amount of time on an ad-hoc basis. Communities collaborate for a longer period with a common interest with which to form a stronger collaborative group. Teams are a formally defined and structured group who collaborate together for a clearly defined amount of time to complete a specific objective.

Numerous classifications exist to structure the different processes that occur between group members. One such classification by Nardi is influenced by the concepts of activity theory [81], which state that individuals conduct activities with an awareness of the surrounding group in which they are located, the rules that restrict their actions and the artefacts available to them. This dictates what each member's regular roles and activities are within a collaborative group and the co-ordination activities required to ensure they link together [11]. These roles and activities are established and re-considered when non-routine situations occur which are addressed through specific collaborative activities. Co-operative processes between group members can be used to determine alternative activities, or if existing activities are not sufficient, co-construction processes determine what new activities are required after reflecting on the organisation of activities and how they meet the goals of the community.

Another classification by Andriessen extends similar concepts to five collaborative processes that are central to group performance: communication, co-operation, co-ordination, information/ knowledge sharing and group maintenance [2]. In this classification communication is an underlying process to all other processes so information can be exchanged effectively and is dependent on common grounding as discussed in section 2.1. Common grounding informs a person's awareness of the scenario in which they are collaborating which has been divided into activity, availability, process, environment and perspective awareness by Steinfield et al. [115].

The remaining processes of Andriessen's classification are either task-orientated processes which are used to complete tasks (co-operation, co-ordination and information/ knowledge sharing) or group-orientated processes which affect the characteristics of a group which can lead to high task performance (group maintenance). Co-operative activities involve two or more group members directly interacting to conduct an activity whilst co-ordination seeks to ensure that activities of individuals or groups are managed to achieve a common goal. Information/ knowledge sharing processes are the ways in which knowledge is spread from individuals to other group members. Group maintenance processes contribute to the amount of trust, cohesion and social identity that exists between group members that results in a effective long-standing group. Trust can affect how group members expect others to conduct their actions and how dependent they can be on these, with cohesion and social identity affecting the norms that the group converges towards according to the common characteristics of its members.

2.2.3 Combined Classifications

There are other classifications that combine elements of the aforementioned classifications to address specific scenarios such as the awareness evaluation model by Neale et al. [82], which combines different elements of team processes and common ground in mainly distributed scenarios. This orders the use of similar team processes to inform a level of work coupling between group members. Work coupling is a measure of how dependent on other's activities each member is and the degree to which they share common or individual goals. Higher levels of work coupling require each member to maintain a better knowledge of the group's context and utilise more frequent and higher quality forms of communication to coordinate actions. Only light-weight interaction is required for low levels of work coupling, increasing through information sharing, coordination and collaboration processes whilst cooperation indicates the highest level of work coupling.

At the lowest level of coupling, group members communicate infrequently and have a low awareness of the actions of others as they largely work towards individual goals. Moderate level work coupling is demonstrated when group members begin to work towards a larger number of common goals whilst individuals maintain their own separate goals. This requires a greater amount of coordination and communication between group members as well as a greater awareness of each individual's context within the group. At the highest level of work coupling, group members are working solely towards common goals with few members having individual goals that are different from these. Many activities are highly interdependent or conducted by the group as a whole. A group will use almost constant communication for each member to maintain a detailed awareness of the group's current context.

2.2.4 System Functionality

Some classifications address the issue in a slightly different way and classify the system functionality rather than the activities of the people within the scenario. The time/space matrix has been combined with the five categories of processes in an attempt to produce a flexible classification for systems by Penichet et al. [88]. Another classification by Borghoff and Schlichter [17] compares systems to previous CSCW systems such as message systems, group editors, electronic meeting rooms, conferencing systems, shared information systems and intelligent agents. By contrast Dix et al. presented a basic cooperative work framework which splits them into three main types, computer mediated communication, meeting and decision support systems and shared applications and artefacts [41].

2.3 Navigation

One long-standing difficulty that users experience in virtual environments over their real world equivalents is that navigation is a non-trivial task [19, 36, 37, 130], forming a potential source of common ground breakdown. However, similar cognitive processes are used to navigate both virtual or real environments (e.g., wayfinding) and can be evaluated through multiple levels of metrics to determine the performance, behaviour and rationale behind one's actions.

2.3.1 Strategies

To navigate any area, humans have to maintain knowledge of where they are within their environment, how this relates to the environment as a whole and how to be able to move from their current location to another. These tasks can be effected by the scale of the environment they are within, and have been categorised into three scales by Weatherford [131]. Model-scale spaces can be viewed in their entirety from a single viewpoint but are too small to move around. Small-scale spaces can be viewed in their entirety from a single viewpoint but are large enough that one can move around within it. Large-scale spaces need to be viewed from multiple positions due to their size and the potential for obstructions such as walls dividing the space.

To navigate large-scale spaces (such as cities) users divide the space into subareas of different scales. One model for this is presented by Lynch where the space is divided into paths, nodes, landmarks, districts and edges [79]. Paths are used to move in a set direction without deviation and are punctuated by nodes where decisions have to be made as to which path to take from them. Districts are large sections of the space which are made of multiple paths and nodes which share some common features which make them identifiable. Landmarks are prominent objects that unlike districts are not entered but are used to identify a person's location and provide context to decisions at nodes. Finally edges serve as a distinguishable border between these different divisions.

These divisions allow the task of going from a starting position to a target point to be achieved through similar strategies regardless of the distance between them. Two main strategies for achieving this are following a route through a series of paired actions at decision points (e.g., go to the church, turn left, continue to the shop and turn right), or moving a specified distance in the direction the point is from the current location (e.g., 500m at a bearing of 120°). These strategies are reliant on route and survey knowledge respectively [121]. These can be complemented through landmarks, optic flow and somatosensory cues from movement (path integration) to help add relative position and distance information to improve their accuracy [78, 110].

Navigation of real-world environments is predominately conducted with full control of method of travel (e.g., walking, driving a car). However, navigational performance can be affected when this control is reduced, e.g., as a passenger in a car [86] or in a pushchair [53].

2.3.2 Evaluation

The complex nature of navigation requires it to be evaluated at multiple levels of granularity to be fully understood. Ruddle and Lessels [100] proposed using three levels of metrics to determine the performance, behaviour and cognitive rationale of users' navigation.

Performance can be determined through metrics such as whether tasks are completed successfully, time taken, distance travelled or number of errors. The metrics chosen must account for the task that is being evaluated in several ways. First, general metrics such as errors are suitable when navigating a restricted environment such as a maze but less suitable for more freeform navigation where time taken or distance travelled would be used.

Secondly, a breadth of metrics should be used to allow a differentiation between structured, effective strategies and less structured behaviour. Users may use strategies that are ineffective when measured with one metric but are highly effective with others. An example of this is when conducting a hub and spoke search strategy which sees backtracking repeatedly to a known point [103] increasing the distance travelled but allowing the task to be completed in less time.

The second level of metrics investigate the behaviours used and may include evaluating the movements around the environment, how different actions occurred during the time or the types of errors made by users. Initial metrics such the total coverage and frequency of visiting subsections of the environment can be combined with viewing visualisations of multiple users' paths through the environment to determine common factors. Such visualisations are covered in greater detail in section 2.4.2.

Classifying the errors that occur can help determine the major difficulties that users are encountering and provide evidence to suggest potential improvements. These can identify whether errors such as users missing targets are due to common issues such as them being outside of their limited view [76], or due to more serious failures in spatial orientation.

2.4 CVE Navigation and aids

A Collaborative Virtual Environment (CVE) is a type of CSCW system that allows multiple distributed users to be virtually co-located and navigate around a 3D environment. This functionality may be required for tasks such as product design [22], urban planning [46] or driving simulation [38, 120]. Collaboration usually occurs through teams conducting tightly coupled processes where task performance is greatly affected by instances of common ground breakdowns.

A variety of different input devices (such as keyboard & mouse, gamepad, props or physical movement) and output devices (flat panel, projector, head mounted displays (HMDs) and augmented reality) could be used to control and view navigation around a VE [20] and some of these options are discussed in section 2.5. However, a high proportion of CVEs are implemented within a desktop based system, resulting in the use of displays with a field of view (FOV) of $\approx$45 degrees and input devices which only allow virtual navigation (the user remains physically stationary).

Although the cognitive processes used to navigate an environment (see section 2.3.1) are similar whether navigating virtual or real environments, people are more likely to become disoriented when navigating VEs than the real world. Key reasons relate to the fidelity of the environment compared with an equivalent real-world setting, the mechanism used for traveling, and a user's field of view (FOV) [76]. These factors cause users to miss spatial and task information [76] resulting in poor navigational performance when compared to real world navigation [93, 117], and requiring additional navigation aids to be provided to assist the user.

Three major forms of aids are additional views, visualisations of previous paths and augmented landmarks, which act in a similar manner to how maps, directions from strangers or GPS systems are used in real life navigational situations. The two main techniques that will be investigated within this thesis are additional views and path visualisations, and these are discussed in sections 2.4.1 and 2.4.2. Additional landmarks are briefly discussed below.

Visually salient and memorable landmarks are used extensively when navigating real world environments but these can be harder to find within a virtual environment due to the restrictions of the display used or the type of data being navigated [24]. Even if a suitable number of landmarks are present within the environment users can have difficulties finding them due to the narrow FOV [76]. Some work by Pierca and Pausch has seen additional or artificial landmarks placed within the environment to address these issues [92].

2.4.1 Additional Views

The types of view that could be provided range along a continuum from human's-eye (purely egocentric) to a north-up global map (purely allocentric). In between are views that are egocentric but tethered to a user's view (e.g., an over-the-shoulder view) as well as forward-up and local maps [132]. The type of map (e.g., forward vs. north-up) that is appropriate depends on the tasks that users perform, but the addition of visual momentum (you-are-here information) to a north-up map produces a solution that is near-optimal for showing the whole of an environment [3]. Visual momentum that show where both oneself and one's collaborators are and can see, should reduce occurrences of common ground breakdown.

Over-the-shoulder views help compensate for the limited FOV provided by most VE systems, by showing more of a user's local surroundings than is visible from a human's-eye perspective and, therefore, addressing FOV difficulties when using desktop displays. Over-the-shoulder views have been implemented in both research systems [69] and computer games [50], and been shown to help users interact [105]. Alternatively, users may be provided with the ability to view an environment from an elevated point, so a larger area is visible even though the FOV itself remains unchanged [34, 42]. However, as the view is brought further back, the user is required to conduct additional navigational actions to view detailed sections of the dataset in sufficient detail.

2.4.2 Path Visualisation

Another family of navigation aids that could be used within a CVE are path visualisations [136]. Path visualisations allow the representation of previous movements within an environment and are mainly used in two ways, to assist users' future navigation of the environment, or to analyse previous navigation to understand users' behaviour.

The use of path visualisations to assist navigation frequently sees them termed trails, taking influence from how marks are made in a real environment (e.g., footprints in snowy or muddy environments show recently and frequently used paths). This type of information is typically not available within CVEs but providing even basic positional information of previous users has been shown to result in more effective search performance by Ruddle [97]. However, users can have difficulties utilising the trails as the quantity of presented trails increases (”trail pollution”) [98].

Analysis of users' movements utilise these visualisations in a variety of ways, from showing a short segment of a single user's movement to investigate particular behaviour [42], to combining large quantities of data from multiple users to understand trends in their movements [25]. These visualisations may be viewed using the same interface and at a similar scale as the original actions [56], or using a different interface and scale [25] depending on the manner in which they are being used.

A path visualisation consists of two main factors, the data being encoded within the path and the way in which this data is visually represented. The manner in which these can be classified is discussed below, with classifications of examples from previous work shown in tables 2.2, 2.4 and 2.3.

2.4.2.1 Data

Path visualisations can encode numerous types of data, from the traversal of a series of documents [39] to the movements of a user in a CVE [36, 56, 97]. Only paths of navigation within a CVE will be considered within this thesis. In a CVE paths can identify the *who*, *what*, *where* and *when* of something that has occurred, similarly to categories used in studies of autobiographical memory [128]. Paths can be encoded with multiple data sources of these elements (see table 2.1).

A CVE allows multiple users to create paths through the environment and identifying these individuals allows knowledge to be associated with them across multiple sessions. These individuals may be investigating the environment on their own or collaborating as part of a group, and group membership may change over time [42].

Although in some CVEs the only action that users will perform is movement around a static environment, others allow for interaction with objects within it [104].

Type	Sources
Who	Individual, Group
What	Movement, Modification, Conversation
Where	Position (2D/ 3D), View
When	Session, Time

Table 2.1: The data types that can be encoded within a path visualisation.

This interaction can consist of adding, modifying or removing objects within the environment, or interactions between users such as conversations [42].

Where a user is within the environment can be recorded through their 2D or 3D position dependent on the data set. Within a 2D data set such as a map, a user's position provides sufficient information to determine what is visible to them. In some 3D CVEs position can still be recorded as their location on a 2D ground plane if movement is restricted to this plane when moving via a walking (or similar) metaphor. A user's view within a 3D environment is also dictated by their view direction, and this can also be recorded to provide a more informed indication of what they can see. Their view can be recorded as their orientation or the position of their view centre in the environment.

The time at which actions occurred can be recorded and represented in different ways. Users will frequently interact over a series of sessions at different times and thus are recorded as separate paths. Within each of these sessions the relative time and order of users' actions can be recorded and encoded.

Where and when data types can be sampled at different resolutions or time frequencies within a path. Time frequencies may vary from constant (every rendering frame), through low frequency but still regular sampling (e.g., every 5 seconds) to on-demand sampling (e.g., users selecting each time they want a sample to be taken). Location data sampling has to be done at a specified resolution, which for position may result in dividing the environment into grid squares of a specified size (e.g., 100m x 100m).

Varying the sampling method can result in the data being used in different ways, for example, sampling position with a high frequency or small grid size could result in a path presenting all the details of a user's movement whilst a low frequency or larger grid sampling sampling could result in some of these details being lost. On-demand sampling might be used to highlight features of the environment that are important to a user as shown by Darken and Sibert [36].

2.4.2.2 Representation

Each of the data types discussed previously can be encoded in a variety of ways by varying basic elements of the visualisation such as the colour, shape, size, orientation and position of symbols. In some cases it may also be appropriate to represent the data through text but this will not be considered within this work. Three examples of visualisation representations will be considered; ball-of-string, markers and heatmaps.

A ball-of-string links a number of positions over a specified time period as a continuous line which provides a basic order for the data but can encode other data through the line's colour and size [97]. These are generally used (see table 2.2) to visualise people's 2D position on a plane within the 3D environment, with the orientation rarely shown (work by Zanbaka et al. [136] being a rare example). Although this representation does not generally encode the frequency with which locations are visited explicitly, users can make simple estimates through the quantity of paths at a location or the width and colour of the line can be varied according to an algorithm which considers the popularity of different paths [98]. The colour of a path can be used to identify individuals or groups [56, 70] and they have been used as both a navigational aid during VE interaction [21, 126] and a tool for demonstrating particular behaviour during post-interaction analysis [25, 136].

Markers can be used to present information at discrete points within the environment, and as they are usually used more selectively than a ball of string representation, many instances encode more types of data within each marker (see table 2.3). The colour and shape of the marker is frequently used to demonstrate the type of action (e.g., by varying the shape from a hand to show manipulation to a footprint to indicate movement [56]) or which individual/group was there [44]. The position of users is also considered in more detail as the 3D position and orientation is more

Ref	Use	Data	Encoding	Other
Borner, 2003 [18]	Post Analysis	2D Position Order Group	Position Order of Line Colour	Also homogeneity of group as size
Chittaro, 2006 [25]	Post Analysis	2D Position Order	Position Order of line	Can be filtered
Dodds, 2009 [44]	Post Analysis	2D Position Order	Position Order of line	
Zanbaka, 2005 [136]	Post Analysis	2D Position View Order	Position of point Orientation Order of line	
Brown, 2009 [21]	Navigation Aid	2D Position Order Frequency	Position Order of line Colour	
Grammenos, 2006 [56]	Navigation Aid	3D Position Order Individual	Position Order of Line Colour	Fades over time
Hoobler, 2004 [70]	Navigation Aid	2D Position Order Group	Position Order of line, Size Colour	
Ruddle, 2005 [97]	Navigation Aid	2D Position Order Time Individual	Position Order of line Colour Colour	
Ruddle, 2008 [98]	Navigation Aid	2D Position Order Frequency	Position Order of line Width/ Colour	Gene paths used to calculate primary paths
Vosinakis, 2011 [126]	Navigation Aid	2D Position Order Group	Position Order of line Colour	Filters to selected semantic

Table 2.2: Examples of paths with a string representation.

Ref	Use	Data	Encoding	Other
Borner, 2003 [18]	Post Analysis	2D Position Modification Time Individual, Group	Position Shape Colour Colour	
Dodds, 2009 [44]	Navigation Aid/ Post Analysis	2D Position Conversation Group	Position Colour Colour	Activiated to replay sound
Darken, 2001 [35]	Navigation Aid	3D Position View	Position Orientation	
Grammenos, 2006 [56]	Navigation Aid	3D Position View Time Modification Individual	Position Orientation Text Shape Text	Fades over time. Can be searched and filtered
Hoobler, 2004 [70]	Navigation Aid	2D Position View Group	Position Orientation + view cone Shape, Colour	
Iaboni, 2009 [71]	Navigation Aid	3D Position Orientation Modification	Position Text Text	

Table 2.3: Examples of paths with a marker representation.

Ref	Use	Data	Encoding	Other
Chittaro, 2006 [25]	Post Analysis	2D Position Frequency	Position Colour	Can be filtered
Chittaro, 2006 [25]	Post Analysis	View Frequency	Position Colour	Of non-navigable locations. Can be filtered
Elvins, 2001 [49]	Post Analysis	2D Position Frequency	Position Colour	
Peponis, 2004 [89]	Post Analysis	View Frequency	Position Colour	
Zanbaka, 2005 [136]	Post Analysis	2D Position Frequency	Position of cell Colour	
Kappe, 2009 [75]	Navigation Aid/ Post Analysis	2D Position Frequency	Position Colour	
Hoobler, 2004 [70]	Navigation Aid	2D Position Order Group Frequency	Position Colour Colour Colour	

Table 2.4: Examples of paths with a heatmap representation.

frequently encoded than with the ball of string representation [56, 70]. The time associated with a marker is rarely explicitly encoded (e.g., through colour [18]) but has been stored as additional information associated with a marker [56].

Heat-maps consider where and when users have been within in the environment (see table 2.4) by varying the intensity of a colour at a position based on how recently a user was there [25]. The 2D position is considered almost exclusively with a couple of examples considering users' view position [25, 89]. They generally represent the aggregated data of many users over a long period of time without explicitly encoding this in a way so that an individual or specific time can be determined. They are almost exclusively used as a post interaction analysis tool, but some work has integrated them within standard VE use, albeit still as more of an analytical than navigational aid [70].

2.5 Display Arrangements

CVEs can be implemented using a variety of display and input devices, ranging from a conventional keyboard, mouse and desktop display, to cave automatic virtual environment (CAVE) [83] set-ups which track a user's body movements and use these as an input device. Other environments immerse the user by utilising head mounted displays (HMDs) to render the virtual environment in an otherwise empty room that the user physically walks around [106].

These input and output devices are used to try and reduce the problems that desktop based systems have in a narrow FOV and reliance on purely virtual navigation. The issues many of these devices have are that they are reliant on cumbersome installation and set-up processes, specialist software or a large financial outlay which makes them undesirable for regular or non-specialist use. As with any maturing technology these issues are being reduced over time and devices such as the Oculus Rift [127] are pushing the mainstream adoption of these devices.

Despite this, desktop systems will still be utilised for the foreseeable future and so only devices that extend this paradigm to a limited extent will be considered within this work. The major examples involve tiled displays with appropriate features of other extremes (tablets and CAVE's) considered. Lastly, the use of combinations of these devices at the same time are investigated.

2.5.1 Multiple and wall-sized displays

The ubiquitous use of computing has resulted in the growing use of multiple displays for everyday tasks, the increased display space allowing for peripheral content to remain visible and quicker context switching [14, 60]. In CVE's the use of double or triple display setups allow for a FOV of $>$120 degrees, resulting in greater spatial awareness [76]. Physical navigation can be provided through low cost tracking systems such as MS Kinect [80].

Using multiple displays to allow for an increased FOV can be further extended to wall-sized displays which provide users with an even larger FOV as well as the ability to physically navigate the data visible on the display. The tiled nature of the displays allows a display area that is similar to projectors but without the loss of pixel density from desktop displays. This allows users to conduct zooming movements to acquire additional detail from the display in addition to sideways movements at a fixed level of detail, as shown by Ball et al. [9].

Conventional input devices such as a mouse and keyboard may be used with these displays but may suffer from issues due to the size of display increasing target acquisition and selection difficulties [96] and restricting their freedom to move around the display whilst interacting. Other input devices that allow users to be untethered (such as Flock of Birds tracked gloves [5]) can be utilised but require modifications of the interaction metaphors used, suffer from inaccurate tracking and require additional training to use.

Implementing CVE's on wall-sized displays has become easier in recent years, but the hardware setup used is still a deciding factor in what software and middleware is utilised. The wall-sized display at Leeds is constructed from 20-inch displays, each with a resolution of 1600 x 1200 pixels, arranged seven displays wide and four displays tall (see figure 1.1a). This provides a total viewing area of 3.06 x 1.33 metres, a resolution of 54 million pixels (11200 x 4800 pixels) and a horizontal field of view (HFOV) of 137 degrees when 0.6 metres away from the display. A cluster of seven machines render output to this display, each of which contain two Quadro 1800 cards that render to two displays each (four per machine).

Running applications on this hardware can be done in a number of ways. Standard game or VR engines such as Unity [119] can be extended to be used on the cluster of machines and wall-sized display, or middleware such as VRJuggler [15], SAGE [73] or CGLX [45] can be used as part of a more bespoke solution. Previously developed software in Leeds utilises VRJuggler to setup and synchronise execution across this cluster for both 2D and 3D datasets with suitable performance [55, 122]. The ability for the same code to be deployed on more traditional single desktop machine and display arrangements through changing a configuration file is another reason to utilise this method of implementation.

2.5.2 Other displays

Other displays that have similar size differences to desktop displays as wall-sized displays include room (CAVE) and mobile (smartphone, tablet etc) displays. CAVE environments provide similar FOV improvements as wall-sized displays alongside more capable interaction devices. This immersive environment has been shown by Steed et al. to provide users with greater spatial awareness of others [113] and by Roberts et al. to allow users to perform better at tasks relying on object manipulation [94].

Smaller form displays often provide a reduced size and resolution display as well as different interaction methods to those used in desktop displays. The reduced display size and resolution has seen the use of off-screen rendering to navigate environments by Ginige et al. [54] whilst the ability for physical movements of the device to be used to control a user's navigation can benefit navigational performance over purely virtual navigation, as shown by Rohs et al. [95].

2.5.3 Heterogeneous Displays

Different displays can be used within both collocated and distributed collaborative scenarios. Collocated collaborations may use different displays to individualise collaborators' interaction or view of the data based on their role requirements. Similar concepts can be applied to distributed scenarios but differences here can also be driven by organisation or logistical constraints on collaborators.

Display heterogeneity has been explored in CVEs and has been shown to affect both high level collaborative behaviour and low level navigational ability for users. The effects on the collaboration vary depending on the functionality provided by the different displays. Users take a more commanding role if they have greater immersion than other users [111] or can manipulate objects within the environment more easily [112]. These differences allow quicker acquisition of task information than less immersed users. Common ground breakdowns increase due to these inequalities unless users correct their natural expectation that other collaborators have the same capabilities as themselves as highlighted by Axelsson [7].

2.6 Scenario

A scenario in which CVEs may be useful is in the analysis of earth sciences data. Organizations involved often have staff spread around the globe, so CVEs could dispense with the need to physically travel to one location to conduct that analysis. However, not everyone in the organization will necessarily have access to the same infrastructure, for instance, one might have a high resolution wall-sized display, but others may have a more modest setup (e.g., desktop) or be based in the field.

2.6.1 Current scenario

Collaboration with Dr Douglas Paton from the University's Centre for Integrated Petroleum Engineering and Geoscience (CiPEG) provided details on existing practice in this area where data is mainly captured in 2D through paper-based artefacts. This data encapsulates scales ranging from outcrops of only a few metres, to entire mountain ranges of several hundred km^2. Analysts initially analyse specific geological features before applying their findings to various scales across the dataset to iteratively improve the geological model of the area. Specialist 3D visualization and processing software (such as Midland Valley's Move [123]), and general tools such as Google Earth, are being used more frequently as they allow the data to be analysed both at and away from the survey location.

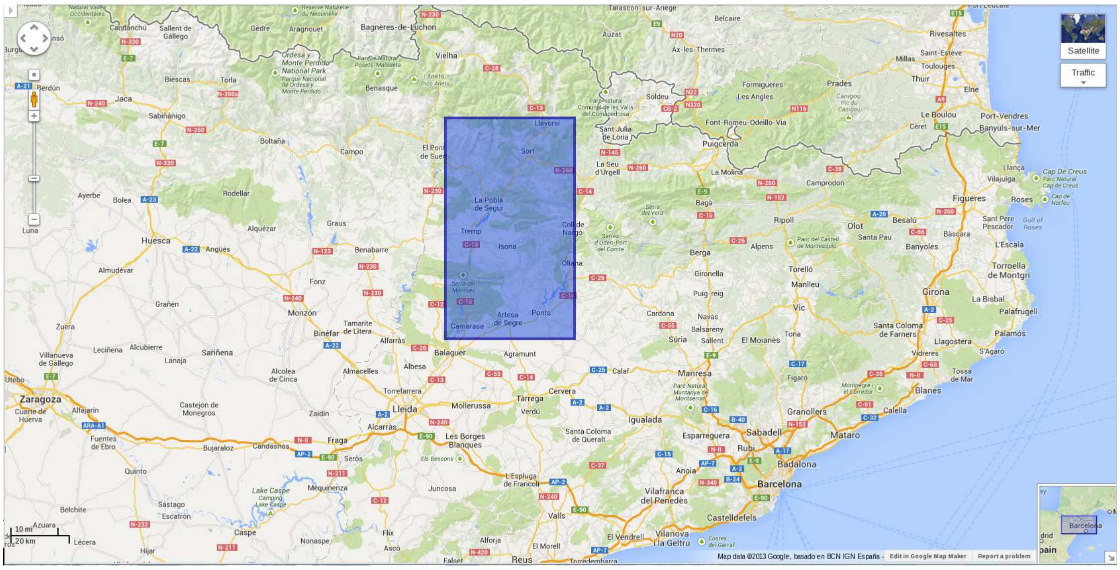

Figure 2.1: The location of the data area used. Map data ©2013 Google, basadoen BCN IGN España.

Current MSc fieldwork visits an area of the Spanish Pyreness over approximately a week, during which multiple outcrops are visited each day for 30 minutes to 2 hours per outcrop. Group sessions are given by the trip leader to the students before the trip to introduce the area and each day's activities, and at the end of each day to recap and update the knowledge acquired to that point. These sessions utilise a mixture of traditional 2D artefacts and Google earth flythroughs.

An area used during one of these days is a 42.5 x 74.6 kilometre area shown in figure 2.1. In this area the southern part is fairly flat, approximately 230 metres above sea level, and contains a number of lakes, rivers and small hills, The northern part of the dataset is mountainous, with a maximum height of 2866 metres, and contains a variety of features that would be analysed when conducting fieldwork in the area.

2.6.2 Use within a wall-sized to desktop display CVE

The dataset presented above is well-suited to being analysed within a CVE on both wall-sized and desktop displays. Analysis of this data requires an accurate and comprehensive understanding of the data in 3D and at different levels of scale. The ability to view this data at different scales and viewpoints within a CVE should be beneficial to users' spatial knowledge acquisition over the 2D-only sources that are currently used.

The use of CVEs within this field is still relatively rare, with both expert and novice analysts likely to be novice navigators within VEs. This means that findings derived from navigational tasks conducted by non-specialised participants are likely to be applicable to many of the real users of this geoscience scenario.

The use of a wall-sized display, with a large FOV and ability to conduct physical navigation, alleviates two major factors of navigational difficulty with desktop VEs. However, desktop displays are still likely to be used extensively in related analysis due to the type of task (e.g., use of specialised software for the creation of the geologic model for a specific outcrop), or restricted availability of wall-sized displays (e.g., using a laptop when out in the field).

Although this scenario could be adapted to allow many forms of collaborative navigation to be used, this thesis will focus on master/ slave navigation. This would allow current material for the MSc fieldwork to be quickly adapted, but would also be used for more general ”Come here! Look at this” tasks that are prevalent when new or updated data is shared between wall-sized and desktop display users. Lastly, This form of collaborative navigation allows consistent stimuli to be provided to multiple users during experiments.

This scenario could potentially be adapted to any location on the time/ space matrix (see section 2.2.1), with the choice of collocated or distributed participants affecting the amount of shared context and synchronous or asynchronous collaboration affecting cost of repair actions. Although collocated activities may occur (e.g., in pre trip briefings), the majority of activities will utilise the system in a distributed manner so this will be used for this scenario.

Synchronous and asynchronous activities could occur at different points of the workflow, with synchronous collaboration placing fewer constraints on grounding and repair actions. This would result in less instances of observable breakdown than in asynchronous collaboration. To ensure that the highest potential level of breakdown is mitigated against, this scenario will utilise asynchronous collaboration.

When navigating this dataset within a CVE, features of the environment can be used as part of the navigational strategies described in section 2.3.1. Features such as rivers, valleys and ridges could be used similarly to how paths, nodes and landmarks are used within city navigation [79]. There are a small number of large, visually salient features such as lakes which can be used as global landmarks, but a large number of geometric features such as hills and valleys are difficult to distinguish from each other, reducing their suitability for navigational use. However, these can be used in combination with other features such as forests, farms and towns to provide a good source of optic flow.

Navigation of this environment requires that the user can freely move in x, y and z directions and rotate via changes in heading and pitch, a more challenging form of travel than when movement can be restricted to a floor plane [20]. Two interaction devices that have been traditionally used with desktop displays, a gamepad and mouse/ keyboard, will be used to provide either identical or similar controls between the two display types (wall-sized vs. desktop).

The use of this multi-scale 3D dataset with few suitable landmarks, freeform travel methods and use of different interaction devices should provide a close to worse-case scenario that will accentuate occurrences of common ground breakdown. However, the use of short tasks will ensure that the actions leading to instances of breakdown can be comprehensively analysed.

2.7 Identifying and addressing Common Ground Breakdown

To reduce the level of breakdown to a workable level the major types of breakdown that occur most frequently in this specific scenario must be identified and addressed. This section presents the major sources of breakdown within the scenario presented above before discussing examples of functionality that have previously been used to address similar breakdowns.

2.7.1 Sources of breakdown

Collaborators within this scenario have to conduct extensive navigation within the CVE, a traditionally difficult task (see section 2.4) which will result in frequent spatial orientation breakdowns. These breakdowns occur when they misunderstand the position of themselves or their collaborator within the environment. When navigating via survey navigation users can make errors in the angle or direction of their movement, whilst errors in landmark identification or direction will leading to failures in route based navigation. These issues can occur at different scales within the environment, meaning small errors navigating between regions can result in great difficulties at the scale of individual outcrops. The ability for the wall-sized display user to to physically navigate results in the desktop display user suffering from an increase in these breakdowns.

Data visibility breakdowns occur when collaborators cannot see the same areas of the environment, usually when collaborators are in different positions through independent navigation. However, the increased visible area of wall-sized displays means its user can see parts of the data that are not visible to a desktop user when they are in the same position, increasing the frequency of these breakdowns.

Spatial orientation and data visibility breakdowns will affect users' navigation and therefore will be measured through analysing navigational performance and behaviour metrics, which were discussed in section 2.3.2.

Another form of breakdown is action visibility where one user does not know what a collaborator is doing. In this scenario, this will frequently be due to navigation of the environment without their collaborator being aware they have moved but could be due to other actions within the environment such as manipulating the data. It can also occur when collaborators are conducting actions that do not require interaction with the CVE. These actions may see them interacting with other data sources (e.g., writing notes) or conducting further investigation of the environment without interaction.

Finally, there will be a number of breakdowns in understanding of the earth science content, similar to the breakdowns in any content common ground. These types of breakdowns will be present in any collaborative scenario and should only be considered when the more frequent breakdowns discussed above have been addressed. As such these breakdowns will not be investigated as part of this thesis.

2.7.2 Reducing breakdown

Some ways to reduce similar breakdowns to those within this scenario have been used within previous examples of collaborative navigation (data visibility and spatial orientation breakdowns) and awareness (action breakdowns).

2.7.2.1 Collaborative Navigation

Many collaborative systems utilise the concept of What You See Is What I See (WYSIWIS), to show the same data and view to all users, resulting in low instances of data visibility breakdowns. However, a strict interpretation of this concept can cause difficulties in coordinating actions across different areas of the environment, resulting in instances of unexpected actions occurring and can be inflexible to differing demands of users. To address this, Stefik et al. suggested allowing WYSIWIS to be relaxed in four main ways; display space (e.g., specific display elements), time of display (e.g., delaying data updates), subgroup population (e.g., providing functionality on a per person basis) and congruence of view (e.g., providing alternative views) [114].

Relaxing the display space allows users to have private workspaces separate from what is displayed to the group at large [77, 109] and not display user interface elements such as the local mouse pointer to the rest of the group. Delaying the time of updates also allows private changes to be reviewed before allowing others to view them, or multiple operations to be combined together. Allowing population relaxation allows different tasks to be competed by different subgroups as well as personalised functionality.

Varying the strictness of view congruence can allow users to have shared or differing views of the data. Each user could move individually from others (equal partners), have their movements automated by the position or movements of others while still having a degree of control themselves (guided navigation), or follow the exact movements of another user with no active control of their own movement (master/slave or yoked navigation). Users may use different forms of collaborative navigation depending on the type of activities being performed. Different views may be required when users wish to conduct loosely coupled activities, while shared views are more suitable when conducting more tightly coupled activities.

Spatial breakdowns are affected by increased movement automation removing the need to conduct low level movement actions, improving navigational performance of the local area, but also resulting in more limited acquisition amount of landmark, survey and route knowledge which can limit future navigation performance [19, 48].

A shared view ensures all users have the same data available to them, even though a level of data visibility breakdown may remain due to misalignments in each user's current focus. These issues are exacerbated with the use of uncoupled views, where each user may have vastly different viewpoints on the data. An extreme form of relaxed view is equal partners where each user can fully control their view of the data. However, users will still frequently require information on what is visible in others' views and this can be achieved through the use of pointers and view cues.

Pointing in collocated settings is an important way to direct the attention of others to items of interest at specific times [10, 134] and can be replicated by displaying the normally local mouse pointer to remote collaborators as a telepointer [57]. Movement of these pointers in specific gestures allows users to convey a large amount of knowledge about the data to others [65]. Another solution is to use representations of what is visible to collaborators (e.g., wireline view frustums [47]) so that users are made aware of when others are viewing interesting data that may not be visible to them.

Additional views which constantly show areas of data visible to others can be used [62], allowing a natural way of repairing visibility breakdowns [4]. Alternative views can also be used, such as peripheral lenses which widen users' FOV [69] to reduce the occurrences of data visibility breakdowns as well as making it easier to maintain awareness of the location of fellow collaborators in the small environment.

However, all of these techniques only inform the user of elements that are not visible within their view and still require them to conduct additional actions to be able to see these areas (or see them without excessive distortion), maintaining a level of common ground breakdown in these scenarios. Additionally, users may have difficulties establishing when this gaze is important, or in determining the target they should be viewing with sufficient accuracy.

Another way that these breakdowns have been addressed is to allow users to teleport to another collaborator's position[43], but these techniques only show what the collaborator can currently see, and doesn't inform them of surrounding and previously visited areas.

2.7.2.2 Activity Awareness

The concept of awareness is frequently discussed in CSCW literature as a general concept, and as instances of designated types such as peripheral, group, activity, workspace awareness [59]. The exact definition and wording of these types can differ, but one major division can be based on the systems providing peripheral or workspace awareness.

Peripheral awareness systems focus on providing users with information on their physical and social context so that they can plan how to initialise different forms of collaborative activity effectively [87]. Examples by Bly et al. make users aware of other collaborator's status through an always-on video/ audio link of a remote location [16], or by Voida at al. open and shut a remote worker's office door depending on their availability [125]. However, systems and aids for these types of awareness will not be extensively considered within this thesis.

Workspace awareness focuses on providing assistance in users' active workspace and their understanding of who is present, as well as the type and location of activities that are occurring [61, 63, 64]. Workspace awareness encompasses information such as the identity of other collaborators, their location within the data, their actions, activity level within the workspace and previous changes they have performed to assist grounding between collaborators. A number of examples of this functionality are provided in the context of distributed software development by Steinmacher et al [116].

Some of this functionality is similar to that provided by peripheral awareness tools, such as the presence of other collaborators and details on the activities they are conducting, but the difference is that this information is focused within the context of the current workspace. An example of this is within the Virtual School system by Carroll et al [23] which provides information about activities that are not within an individual's or the group's current focus, but still within the scope of the current shared workspace.

The information that is provided can be used to personalise system functionality and data representation to remove unnecessary user interface elements or create different interfaces based on unique tasks their role requires. An example by Convertino et al allows users such as a weapons expert to conduct specific tasks (e.g., marking out a minefield), and present a summary of this data to others [30]. The explicit representation of a user's role also allows other users to make implicit assumptions of a user's characteristics, reducing the need for explicit grounding behaviour in time critical scenarios such as emergency planning [33].

The recording and exposition of previous interactions is also considered as part of edit and read wear tools introduced by Hill et al. [68] which record and visualise each user's actions within basic documents and web pages [1], as well as larger and more complex visual spaces such as maps [51]. This interaction data can exposed to other users within looser coupled and/ or less rigidly defined groups through social navigation functionality. An example of this is presented by Wong et al. and highlights potentially interesting areas where previous users have been to help a new user to begin exploring the data [135].

2.8 Evaluation Methods

To investigate and evaluate the effect that potential aids may have, a suitable method and choice of activities needs to be decided upon. Potential methods that could be utilised include log analysis, observation, controlled laboratory study and interviews. Methods such as observation and log analysis seek to discover the types of behaviour that occur naturally in realistic scenarios, but their unstructured nature can lead to difficulties in finding evidence of specific phenomenon and to account for the large number of variables present. More controlled laboratory studies can vary specific factors over a smaller subset of tasks to find quantitative data with a lower number of participants. Additionally combinations of different methods may be used to address the weaknesses in one or to give both quantitative and qualitative data.

The experiments in this thesis were conducted as controlled laboratory studies to analyse the effect of specific factors. Logs of participants actions were also taken so that the behaviour in specific situations could be analysed and categorised.

The type and order of activities that participants perform must be carefully considered. Due to the fairly limited period of time that each participant will conduct the experiment (a few hours at most), it is important to ensure that they are given sufficient time to become proficient with both the interface [101] and the task [129] that they will conduct to ensure that the findings are representative of trained performance. An effective way of ensuring this is to use a three stage method where participants first practise using the interface while completing a simplistic task, then practise the task that they are required to perform, before finally conducting the test stage [106].

2.9 Summary

This chapter has introduced the areas relevant to the research that is presented in the rest of thesis. Common ground and its associated processes were initially introduced before considering how larger collaborative contexts can be classified through CSCW. Navigational strategies and evaluation methods are discussed before introducing aids that are used to alleviate the difficulties of navigation within CVEs.

Following from this, different forms of display that can be used to view CVEs were discussed, concentrating on multiple and wall-sized displays, before considering the issues of heterogeneous display use. Lastly, a scenario in which common ground breakdown within wall-sized to desktop CVE use can be explored is presented before identifying the types of breakdown that will occur. Lastly, the evaluation methods that will be used are discussed.

Chapter 3

An experimental collaborative platform

3.1 Introduction

This section will introduce the CVE platform that was developed to investigate the issues of common ground breakdowns within the scenario described in section 2.6. The requirements for this platform will be outlined first before discussing the format of the data that will be processed by the resultant system. Following from this the hardware, structure of the developed software and format of the output logs used in the final system are described.

3.2 Requirements

The developed platform had the following requirements:

- Allow users to navigate a 3D geographical dataset with navigational aids, graphics and control mechanisms that are representative of the current state of the art.

- The platform must be controllable with standard desktop control input (mouse/ keyboard and gamepad).

- The rendering output must be viewable over a variety of display arrangements including desktop and wall-sized displays.
- Provide flexibility for changes such as the implementation of navigational aids which modify processing of control input, multiple views of the data or the addition of other visual aids within the dataset.
- Record and output the movements of users as a log file which can be read later by this or other software. The software needs to be able to read these logs later to recreate the movements used, represent the paths for later analysis or to calculate statistics.

3.3 Data

The dataset that will be utilised through these studies corresponds to the area described in section 2.6.1. The Institut Cartogrfic de Catalunya (ICC) provides access to topographic data and orthophotos of the Pyrenees mountain range within the Catalunya region. The region mapped in this manner is split into regularly sized sheets of approximately 14 x 9.5 kilometre each.

For each sheet a dxf and sid file are available which provide topographic details and an orthophoto of the area respectively. The dxf files provided the data for a Digital Elevation Model (DEM) consisting of x,y positions within the environment at 10m height intervals. An automated mesh generation algorithm could be used on this data to generate a polygonal mesh but the quality of the mesh would vary depending on the scale of elevation change and tests such as mesh intersection and finding the nearest vertex to a point would be more complex and time consuming than with a regular grid data structure.

Therefore the data for each sheet was re-sampled by dividing each sheet into a 100 x 100 cell grid, averaging the heights of all points within each cell and applying the resultant height onto a single point in the middle of the cell. These points were then saved as a text file with a point per line indicated through its x,y and z position producing a dataset with a total of 24,000 vertices (30,000 x 80,000) which occupies a total of approximately 8.5MB on disk.

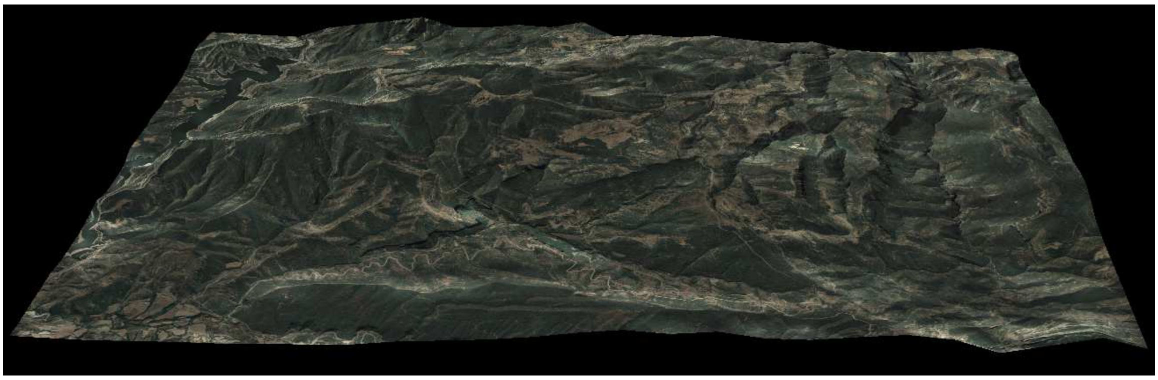

Figure 3.1: Example of the terrain mesh after being texture mapped with an orthophoto.

The orthophoto within the sid file for each sheet was converted into a TIFF image to be overlaid onto the resultant mesh to produce the final terrain. Each TIFF used is approximately 1400 x 950 pixels and occupies a total of approx 100MB on disk, an example sheet of which is shown in figure 3.1. When the sheets of the desired area are combined they produce the environment shown in figure 3.2.

3.4 Hardware

Two hardware set-ups were used, a cluster based system with a wall-sized display output (see figure 1.1a) and a more conventional single desktop system and display (see figure 1.1b). The wall-sized system utilised 7 nodes, each equipped with a Intel Xeon E5520 quad core CPU running at 2.2 GHz, 12 GB of RAM, two Nvidea Quadro 1800 graphics cards and a 1Gb/s ethernet connection. The display was constructed from 20-inch displays, each with a resolution of 1600 x 1200 pixels, arranged seven displays wide and four displays tall. This provides a total viewing area of 3.06 x 1.33 metres, a resolution of 54 million pixels (11200 x 4800 pixels) and a horizontal field of view (HFOV) of 137 degrees when 0.6 metres away from the display.

The desktop system utilised a Intel Xeon E5620 quad core CPU running at 2.4 GHz, 6 GB of RAM, a single Quadro 1800 graphics card and a 100Mb/s ethernet connection. The display was a 30-inch display with a resolution of 2560x1600 pixels and a HFOV of 60 degrees when the participant was 0.6 metres away from the

Figure 3.2: The dataset used throughout the experiments.

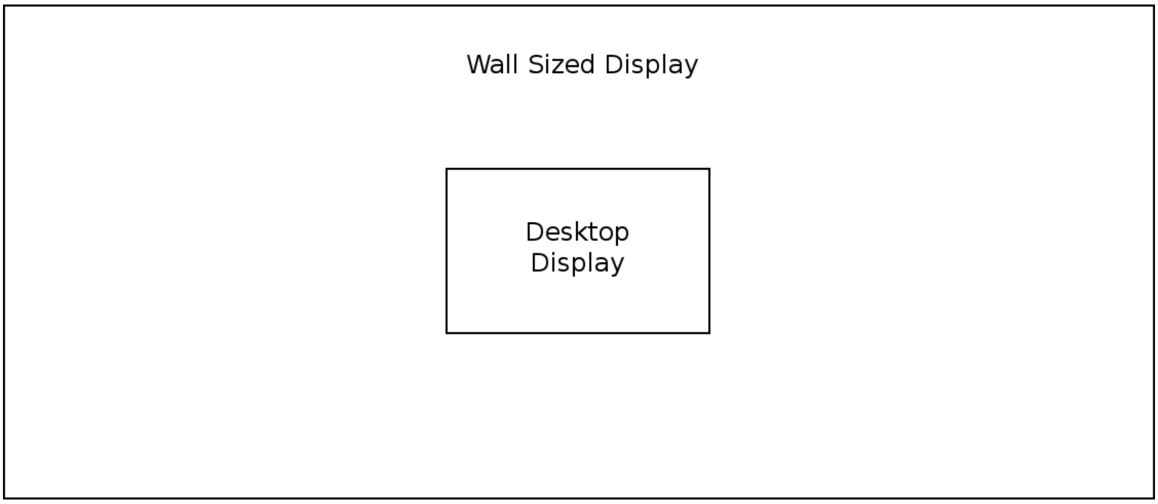

Figure 3.3: The relative size and position of the desktop display relative to the wall-sized display.

display. The size and location of this display relative to the wall-sized display is shown in figure 3.3.

3.5 Software

As discussed in section 2.5.1 the required functionality could be achieved through different VR toolkits, middleware or bespoke solutions. A mainly bespoke solution that utilised VRJuggler was decided upon to allow functionality to be flexibly and selectively implemented depending on the requirements for each experiment. VRJuggler allows different display and control input configurations to be described through XML based Jconf configuration files which are read and processed by a kernel program to initialise the required setup before running a modified OpenGL pipeline. This kernel provides inherited methods at various stages of the rendering pipeline to allow user code to be appropriately executed.

The system consists of 5 main classes, SpainApp, Mesh, Path, Trials and Recorder. The SpainApp class implemented the main VRJuggler kernel and deals with initialisation of the other classes processing the command line arguments, tracks the general system state and processes input device changes. Mesh represents a sheet

Figure 3.4: Example of the path visualisations used to analyse movements.

of the mesh and provides functionality to read the mesh file, render the mesh and processes intersection and nearest vertex tests. The Path class represents a single path through the environment with functionality to read a path file, replay the path over time, calculate statistics based on the path and visualise the path (see figure 3.4).

The Trials class structures a combination of paths so they can be used for the experiments, providing functionality to read the files containing the target locations associated with each path, an interface to access data of the current path, rendering of multiple paths and calculate statistics that are dependent on the target locations. Lastly the Recorder class controls the recording of data to logs, providing methods to start the writing to a log, write data per frame, record actions such as confirming a target and finishing the log recording. The format of these logs are described in section 3.7.

3.6 Interface

Input from the user was given via either a keyboard and mouse (see figure 3.5) or gamepad (see figure 3.6). Mouse movement allowed the user to translate and rotate their view using displacement-based control. Holding the left mouse button allowed translation in a plane perpendicular to the current view direction (at a rate of two metres per pixel), holding the middle mouse button allowed translation in a plane parallel with the view direction (at a rate of ten metres per pixel) and holding the

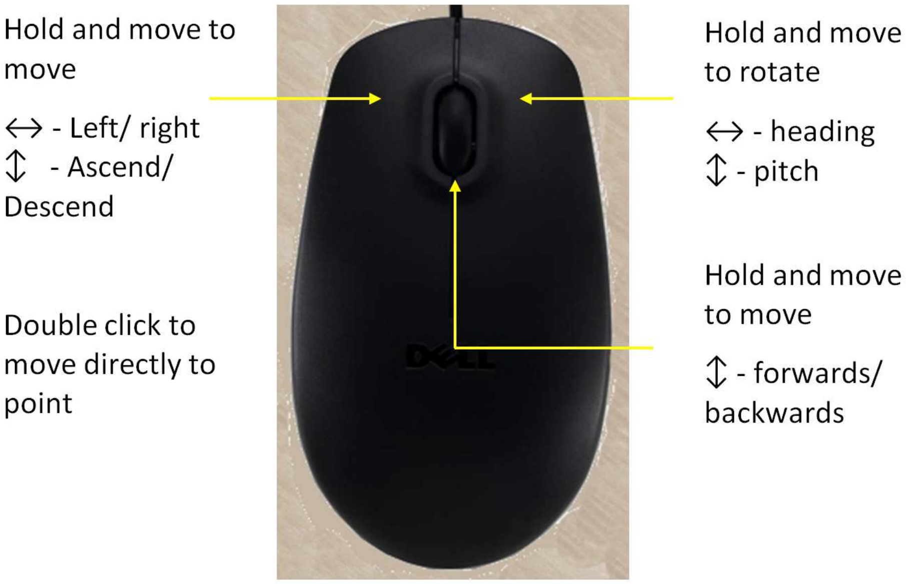

Figure 3.5: Mouse controls.

right mouse button allowed variation of pitch and heading (at a rate of one degree per 20 pixels). Users could move discretely by double clicking with the left mouse button on a point of the landscape. This moved them to a point 3000 metres (or a quarter of the current distance from the selected point if <12000 metres away at the beginning of the movement) from the selected point in one second at a constant speed during which their heading remained constant.

Gamepad controls were implemented via a wireless Logitech RumblePad 2 gamepad, allowing navigation using a flying metaphor with 5 degrees of freedom (translation in X, Y and Z and rotation for heading and pitch). The left joystick allowed movement forwards, backwards and diagonally in the current heading and pitch while two buttons on the back of the gamepad allowed ascent and descent perpendicular to the current view direction. The right joystick was used to vary their heading freely and the pitch between +/- 80°. Translational movements were limited to a maximum of 1200 m/s and the view direction could be changed up to 30 °/s.

With both control methods, collision detection with the ground was implemented so that participants could not pass through the terrain to avoid unnatural movements for the user [72]. Collision detection involved a check to ensure that the user was a minimum distance (50m) above the greatest height of the nearest vertex of the

Figure 3.6: Gamepad controls.

mesh and the 4 vertices surrounding this vertex.

3.7 Path Logging

The logs have to be easily readable by other software and be sufficiently documented so they can be easily understood by anyone else wanting to read them or use them as input into their software. To satisfy this the log file is split into two main sections, a header which describes the structure of the log in a human readable format, and then a body with lines of comma separated values (CSVs) which can be easily parsed by other software. An example of this structure is shown below.

```
* Time, Position (x,y,z), Quaternion (w,x,y,z)
0.000,-19386.199,1022.810,-1085.770,-0.20734,0.04540,0.95459,0.20905
0.017,-19386.201,1022.810,-1085.769,-0.20734,0.04540,0.95459,0.20905
0.034,-19386.203,1022.810,-1085.768,-0.20734,0.04540,0.95459,0.20905
0.050,-19386.205,1022.810,-1085.767,-0.20734,0.04540,0.95459,0.20905
0.067,-19386.207,1022.810,-1085.767,-0.20734,0.04540,0.95459,0.20905
.
.
```

Header information is indicated through the use of a $*$ symbol and describes the order and types of data listed in the log. Following this the position and orientation of the user is initially recorded at a per frame resolution to achieve the highest resolution possible from the initial recording. The (relative) time is also recorded to ensure that the movements are played back over the same time period as the original recording, independent of the rendering frame rate of the hardware being used.

This basic structure was extended for the purposes of recording multiple independent trials required in Experiments 2, 3 and 4. An example of this modified structure is shown below.

```
* # Format Per Trial
* # Trial Number, Trial Type (Visible, Movement Type, Movement Style,
Feature), Trial Distance
```

```
* # Start Trial n #
* Time, Position (x,y,z), Quaternion (w,x,y,z)
* # End Trial n #
* # Trial Number, Trial Type (Visible, Movement Type, Movement Style,
Feature), Time Taken (guided, navigated, difference), Distance Travelled
(guided, navigated, difference), Root Mean Squared, Number of Targets,
Minimum Distance From target, Time to Hit Target,
[Minimum Distance From target, Time to Hit Target]

# 3,Visible,Translate,Follow,Path,46412.14
# Start Trial 1
0.000,-19386.199,1022.810,-1085.770,-0.20734,0.04540,0.95459,0.20905
0.017,-19386.201,1022.810,-1085.769,-0.20734,0.04540,0.95459,0.20905
0.034,-19386.203,1022.810,-1085.768,-0.20734,0.04540,0.95459,0.20905
0.050,-19386.205,1022.810,-1085.767,-0.20734,0.04540,0.95459,0.20905
0.067,-19386.207,1022.810,-1085.767,-0.20734,0.04540,0.95459,0.20905
.
.
42.129,-3073.697,445.072,13055.615,0.33802,-0.12233,0.87745,0.31755
42.146,-3073.697,445.075,13055.617,0.33802,-0.12233,0.87745,0.31755
# End Trial 1 #
3,Visible,Translate,Follow,Path,19.883,42.146,22.263,46412.14,60694.02,
14281.88,1352.1383011392,3,491.64,11.292,1411.56,42.129,465.48,18.693
```

The header again describes the format the file will take but the CSVs values for each trial are pre-ambled by two lines (denoted by a #). The first of these lines contains basic information about the trial and the second indicates the start of the trial. The end of the CSVs is also followed by a line starting with a # to show the end of the recorded data for that trial, allowing an individual trial to be easily found and processed. Another CSV line follows the end of each trial and contains a summary of some of the data that was calculated during execution allowing this information to be easily stripped from the log with a bash script.

3.8 Summary

This chapter has introduced the hardware and software architecture of the system that is used throughout the experiments in future chapters. The manner in which users' paths are recorded, replayed and analysed is described as well as the form of the data that they are navigating. The interface devices used and form of the log file produced by the software is also detailed. The system provides a common platform for the future experiments but could also form the basis of a more general VE across different display setups. This is discussed in section 7.1.5.

Chapter 4

Investigating the causes and severity of common ground breakdown

4.1 Introduction

The theoretical effects of common ground breakdown in this wall-sized to desktop display collaboration have been discussed in chapter 2 but an accurate classification of the frequency, severity and conditions of these breakdown is needed. This chapter describes two experiments with this aim. The first required participants to perform a restricted task of navigating to a single target to ensure that the core task is suitable for further investigation. The second expanded upon this to consider a task more representative of those that occur in real collaborative activities.

4.2 Experiment 1: Constrained navigation to singular targets

An initial exploratory experiment was conducted with the following aims:

- Assess the base level of common ground breakdown that occurs during simple collaborative tasks between wall-sized and desktop users using the platform presented in section 3.

- Investigate the effect of a representative range of standard navigation factors within this scenario to guide the design of future aids design and experiments.

Master-slave collaboration was used, with the experimenter acting as a master and using the wall-sized display, whilst participants acted as the slave and used the desktop display. The participants' task was to observe where the master took them (from a start point to a target; for consistency these navigational movements were pre-recorded), and then navigate themselves from the start to the target.

We therefore had the following hypotheses:

H1.1 Aids that increase the view similarity between desktop and wall-sized displays (e.g., providing a wide FOV context view to the desktop participant) will decrease common ground breakdown between the two.

H1.2 Scenarios that require participants to conduct more complex navigational tasks (e.g., movement involving translation and rotation or where the target is not visible to both desktop and wall-sized participants), are more susceptible to common ground breakdown.

4.2.1 Method

4.2.1.1 Factors

To establish what areas of this scenario required more detailed investigation in later experiments several different within-participants factors were chosen. These were;

- Type of view provided (base vs. context)
- Target visibility on the wall-sized display from the starting position (visible vs. not visible)
- Type of master's path movement (translate vs. translate & rotate)

The first factor investigates what effect inequalities in environmental visibility have, as this is one of the main contributing factors to improved navigational performance with wall-sized displays (see section 2.5.1). Providing the context view (see figure 4.1b) eliminates the inequality in display area and can be used to consider what effect the other features of a wall-sized display (such as physical navigation) may have. This factor corresponds to the data required to test H1.1.

The final two factors investigate if the complexity of the master's movement affects the level of breakdown in the scenario, a phenomenon noted in previous work on guided navigation [19]. The addition of rotation to the master's movement will present participants with more complex movement than the typically straight (translation only) paths that are frequently used [99]. Paths where the target is not initially visible on the wall-sized display forces a wall-sized display user to conduct navigational actions more similar to the desktop display user who can never see the target from the start point. These factors provide data to test H1.2.

4.2.1.2 Measures

The following measures are taken during each trial:

- Success/ failure in finding the target
- Time elapsed during the trial
- Distance travelled during the trial
- Position and orientation per frame

The first three measures allow performance to be assessed while the last measure allows participants' navigational behaviour to be investigated. Instances of common ground breakdown are observable through reduced navigational performance and/ or differences in navigational behaviour.

4.2.1.3 Participants

Twelve participants (six females and six males, mean age = 27.6, SD = 5.9) took part, all of whom (and those in future experiments) were students or staff at the University. The order in which participants used each view (base vs. context view) was counterbalanced. Two separate sets of start/target pairs were used, with half of

(b)

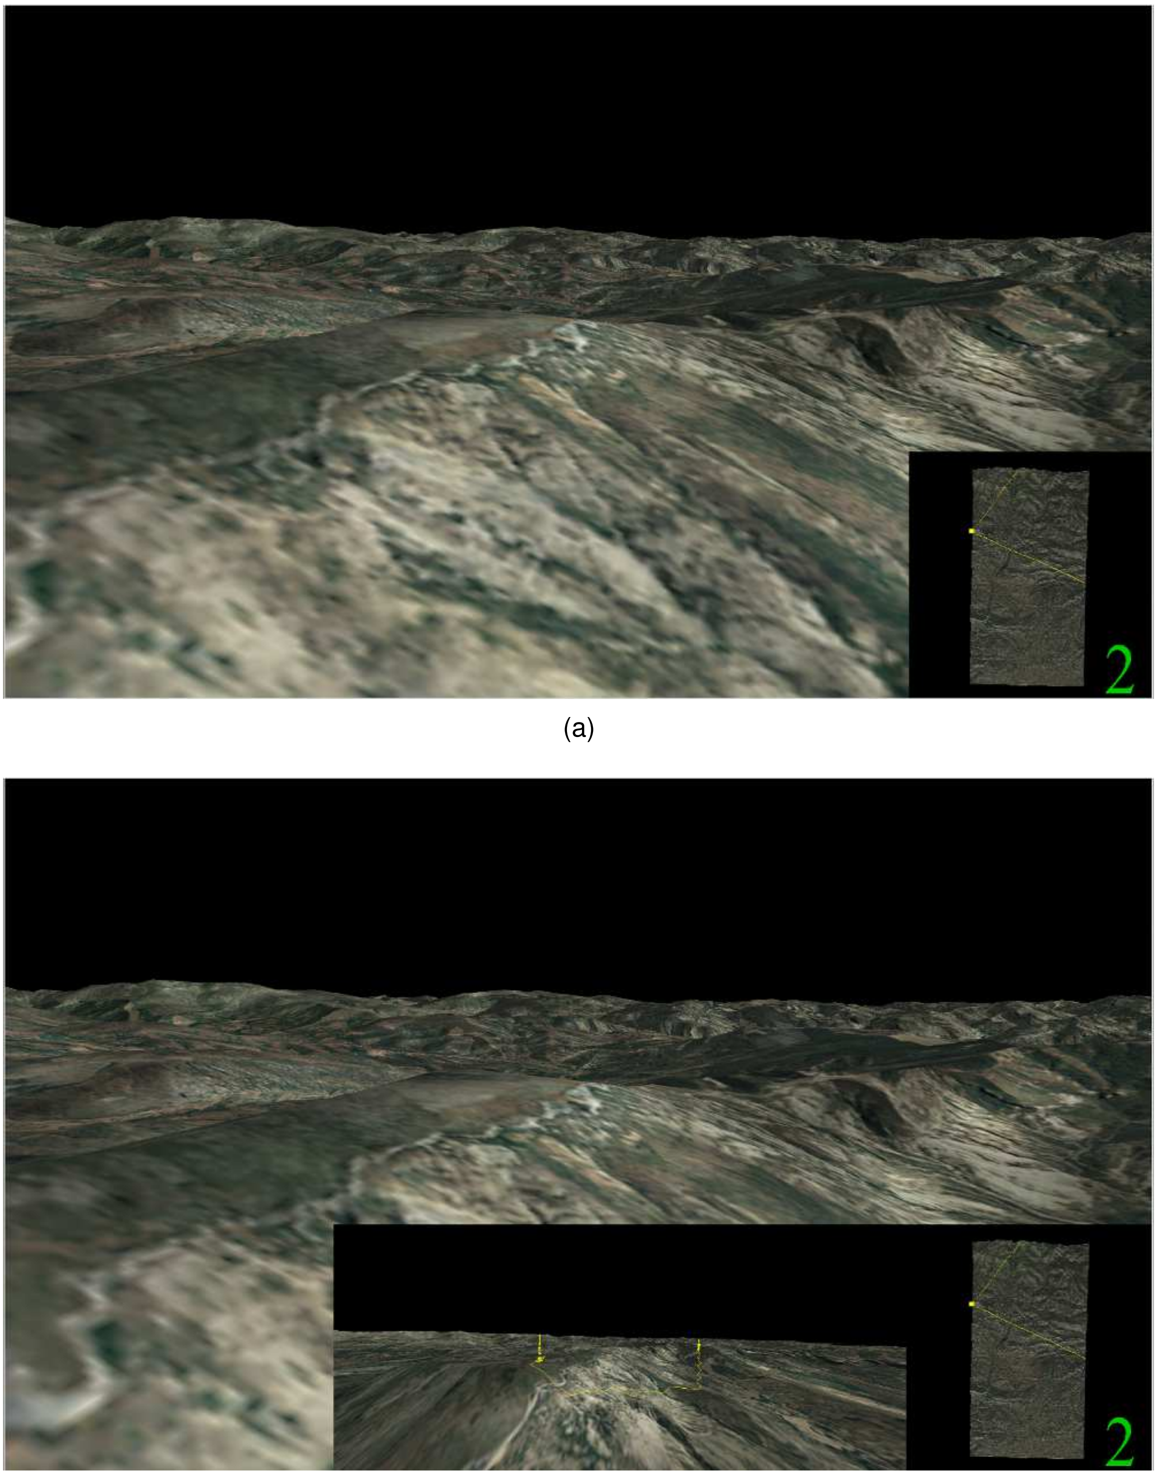

(a)

(b)

Figure 4.1: The view conditions for Experiment 1: (a) Base view showing the global map in the bottom right of the display and (b) Context view along the bottom of the display.

Visibility	Movement	Path Length (m)
Not Visible (8)	Translate (4)	Mean 12792 SD 4762
	Rotate (4)	Mean 14199 SD 2998
Visible (8)	Translate (4)	Mean 12158 SD 2595
	Rotate (4)	Mean 14775 SD 4430

Table 4.1: Split of 16 paths for each block.

the participants using each set with each view, again counterbalanced. All the experiments described in this thesis were approved by our University Research Ethics Committee. All the participants gave their informed consent and were paid an honorarium for their participation.

4.2.1.4 Materials

The input and output devices, controls and dataset that formed the environment used in this experiment were described in chapter 3. Both wall-sized and and desktop users only utilised the gamepad to navigate the environment. The context view utilised by the desktop participants (see figure 4.1b) occupied a 1280x533 pixel space at the bottom of the display and provided a view of the environment with the same 137 degrees FOV as the wall-sized display.

4.2.1.5 Procedure - Master

The experimenter pre-recorded 64 paths, consisting of movements from a start point to the required target, using the wall-sized display. Pre-recorded paths rather than live navigation was used to ensure that each slave participant was subjected to exactly the same movements. The recorded paths were intended to simulate a situation where a user is familiar with a dataset and is interested in particular features within it. The user's navigation is therefore expected to be smooth and direct, which is less likely to cause common ground breakdown than during data exploration, when navigation is more varied.

The 64 paths were divided into four blocks, two of which were used for training trials and two for test trials. The characteristics of the trials in each block are shown in Table 4.1. For half the trials the target was visible on the wall-sized display from the start point, and for the other half of the trials it was not. For half of each visibility the movement was predominantly translation, and for the other half it was translation & rotation.

4.2.1.6 Procedure - Slave

As discussed in section 2.8 the procedure was divided into three stages. First, a participant practised using the gamepad interface to navigate around the landscape. For this, they spent approximately 5 minutes navigating between six clearly-visible targets with the base view, and then spent a similar amount of time navigating between six other clearly-visible targets with the context view.

Stages 2 and 3 had identical formats, with a participant using one of the views (e.g., base view) in the former, and the other view (e.g., context view) in the latter. The format involved two blocks of 16 trials, separated by a five minute rest. In each stage, the first block was treated as training and the second block was the test. The trials in each block were presented in a random order, with eight of the trials involving predominantly translational movement and eight involving translation and rotation. In each eight, four were trials where the target was visible from the start point on the wall-sized display, and four were trials where the destination point was not visible.

The procedure for each trial was as follows. First, a participant pressed a button on the gamepad and was placed at the trial's start point. Then the participant pressed the button again, and was navigated to a target location along a pre-recorded path. Then the participant pressed the button for a third time, which placed them back at the start point, and they navigated themselves to the target. The trial stopped when the participant was within 150 metres of the target or they had taken longer than the cut off time of 1½ minutes (double the time of the longest trial). The targets were indicated by cubes with a letter (see figure 4.2) and texture mapped

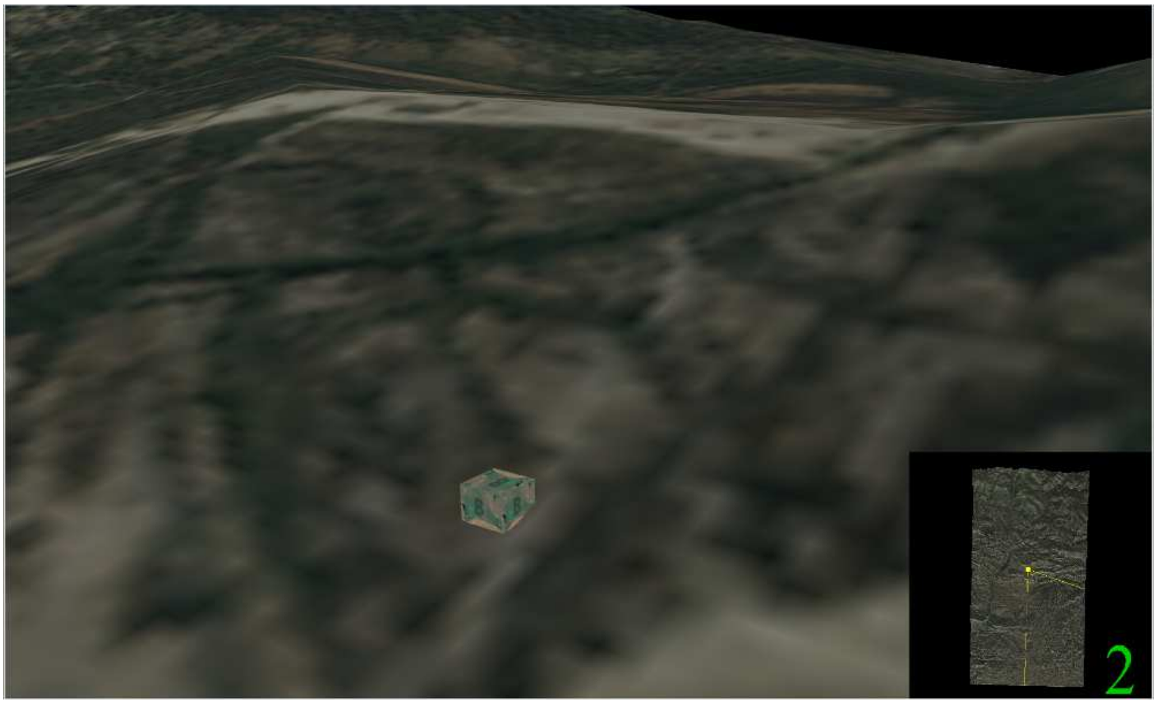

Figure 4.2: A sample target within the participant's view.

similarly to the landscape surrounding it, requiring participants to be close to the target (typically <3000m) to see it. At all times, all the targets of the current stage were displayed (with the same texture mapping as the target but without a letter), to encourage participants to navigate using information memorised during the pre-recorded movements and not to zoom out to try to detect a single target from a distance.

The system used within this experiment (and subsequent experiments) was initially informally shown to other members of the research group to ensure aspects such as the control interface and movement speed was appropriate. The full experimental procedure was subsequently piloted with two participants (the results of which are not reported) to ensure these tasks were of appropriate difficulty and length. Feedback from these methods and observations by the experimenter prompted small changes and bug fixes to the system used.

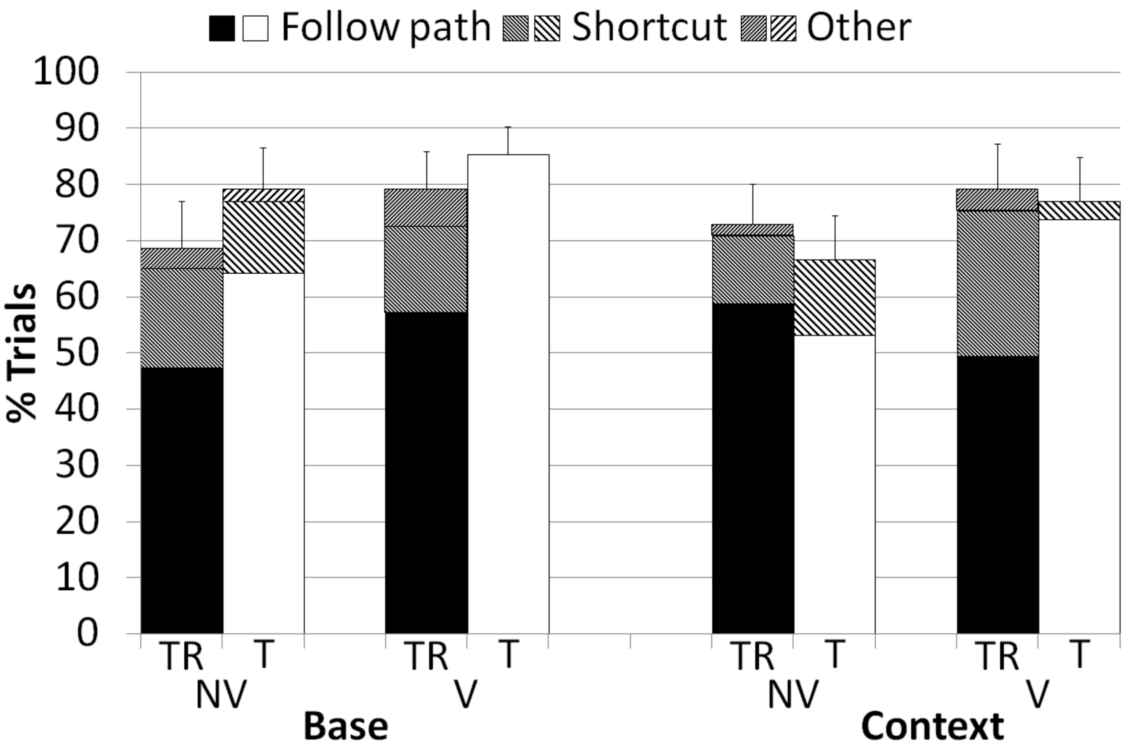

Figure 4.3: Percentage of trials that were successful for each combination of views, target visibility (visible V, not visible NV) and movement type (translate & rotate TR, translation T). Error bars show standard error of the mean for each column as a whole. Each column is subdivided to show the navigation types used in the trials.

4.2.2 Results

The results are divided into two parts. First, the performance of participants is analysed. Second, the navigational behaviour of participants is investigated to explain differences in performance.

4.2.2.1 Performance

The percentage of successful trials was analysed using a repeated measures analysis of variance (ANOVA), which treated the type of view (base view vs. context view), movement type (translate vs. translate & rotate) and initial visibility of the target's position on the master's display (visible vs. not visible) as within-participants factors. Only significant effects (p $<$.05) are reported.

Participants successfully found targets significantly more often when the target position had been visible from the starting point on the wall-sized display (*M* = 80%) than when the target position was not visible (*M* = 72%), $F(1, 11) = 5.25$, $p < .05$. There was also an interaction between view and movement type, $F(1, 11) = 5.08$, $p < .05$. In the base view condition participants were more successful when movement was translate than translate & rotate (*M* = 82% vs. 74%), but in the context view condition they were more successful when movement was translate & rotate (*M* = 76% vs. 72%) (see figure 4.3).

No other measurements (distance travelled, raw or normalized differences between original and travelled distance or time) had significant differences, due to the large range of utilised strategies and effectiveness in finding the targets.

4.2.2.2 Behaviour

It become noticeable that in a significant proportion of trials, participants did not attempt to follow the path that they were guided on, instead some took their own, direct path to where they thought the target was (shortcut). This behaviour has been previously investigated as part of navigation within both real [12] and virtual [52] environments and is similar to the triangle competition task that is frequently used in spatial skill studies [26, 118]. Other trials showed participants taking longer to reach the target than the guided path whilst still trying to follow the original path (getting lost). Some trials also showed that participants seemed to struggle to find the actual target after getting to the approximate area relatively quickly.

Because of these different behaviours each participant's navigation was split into two distinct phases. Firstly, the participant's (attempted) navigation to the approximate target area, and secondly, their search of this approximate area for the actual target. To classify the first stage the original and participant's path were projected onto the ground plane and a vector was taken from the starting point of the trial to the point at which the participant was 1000m away from the start. If this vector was within 45° of the initial direction of the original path then the trial was classified as following the path (see figure 4.4 and 4.5). If participants did not follow the path then their movement was classified as taking a short-cut if they set off within 45° of the

Figure 4.4: Example of a follow navigation type shown as a string path sampling every frame (see section 6.2.1). The original path is shown as progressing from red to green, the participant's path goes from blue to yellow and the target point is highlighted in white. The participant originally sets off in the direction of the path and stays within 1000m of the original path until 3000m away from the target.

vector towards the target (see figure 4.6), other if a participant's path direction was outside this angle.

The percentage of successful trials for which each path type was used was analysed using a repeated measures ANOVA that had four within-participants factors: type of view, movement type, initial visibility, and navigation method (follow path vs. shortcut vs. other). Only significant effects, adjusted for bonferroni corrections (p $<$.025), are reported.

Participants attempted to follow the original path (*M* = 79%) significantly more often than taking a shortcut (*M* = 17%) or any other strategy (*M* = 3%) $F(2, 9) = 280.99$, $p <.01$ (see figure 4.3). There was a significant interaction between navigation method and movement type $F(2, 9) = 5.84$, $p <.025$. Participants attempted to follow the original path more when the path was translate only (*M* = 89%) than when both translation & rotation occurred (*M* = 70%) and a short-cut was used more often with translate & rotate paths than predominately translate paths (M = 24% vs. 10%).

The use of shortcuts during some translate & rotate paths indicates that participants gained a high level of common ground in some scenarios. However this cannot be investigated accurately as they may have a similarly high common ground in

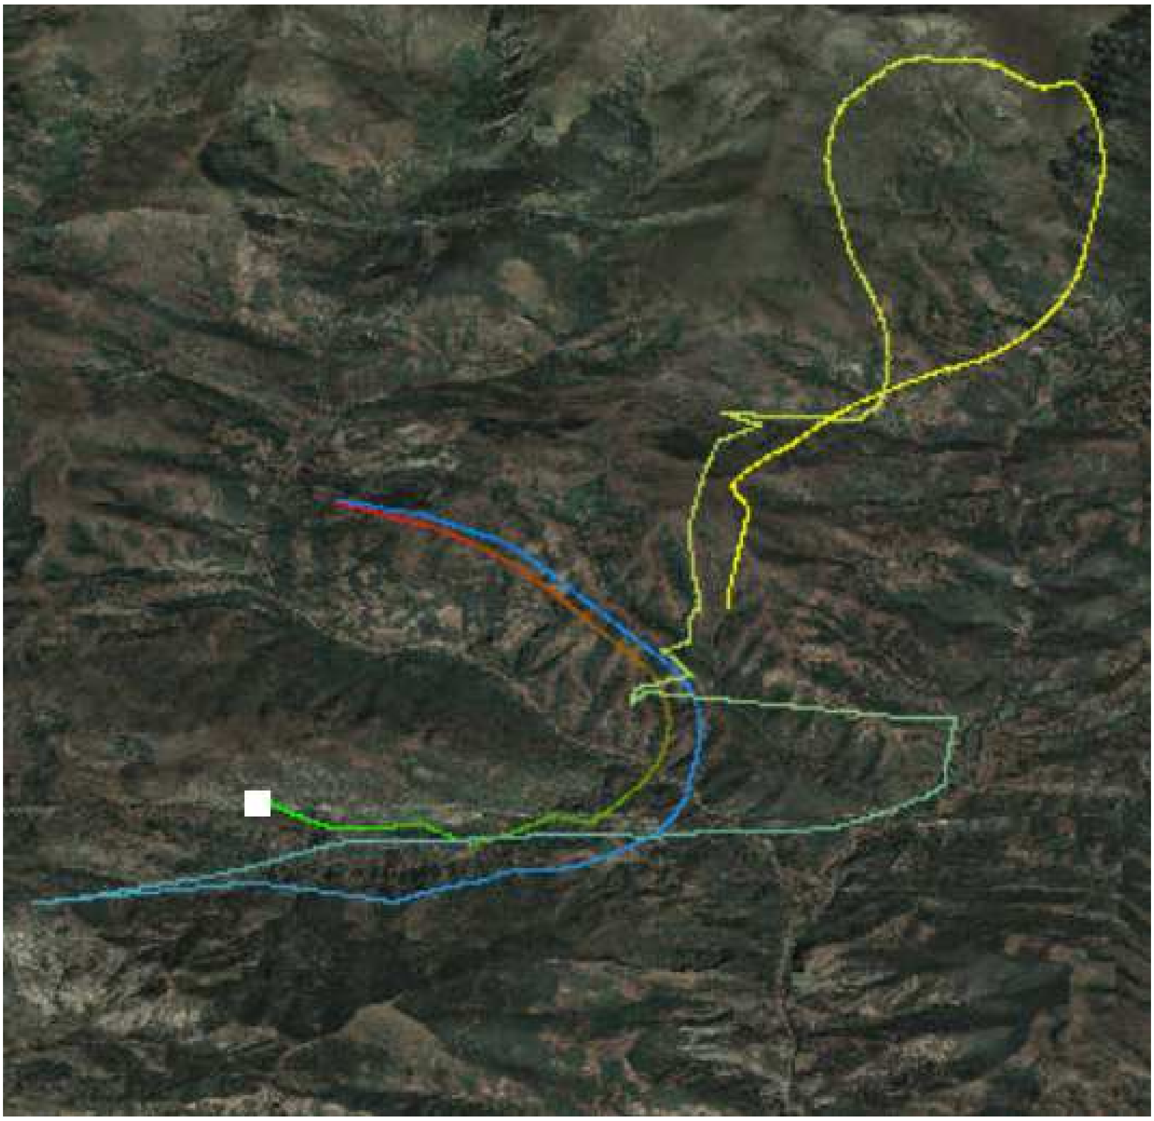

Figure 4.5: Example of a error navigation type shown as a string path sampling every frame (see section 6.2.1). The original path is shown as progressing from red to green, the participant's path goes from blue to yellow and the target point is highlighted in white. The participant starts by following the original path but then turns off.

Figure 4.6: Example of a shortcut navigation type shown as a string path sampling every frame (see section 6.2.1). The original path is shown as progressing from red to green, the participant's path goes from blue to yellow and the target point is highlighted in white. The participant does not set off in the original path direction, instead heading towards the target.

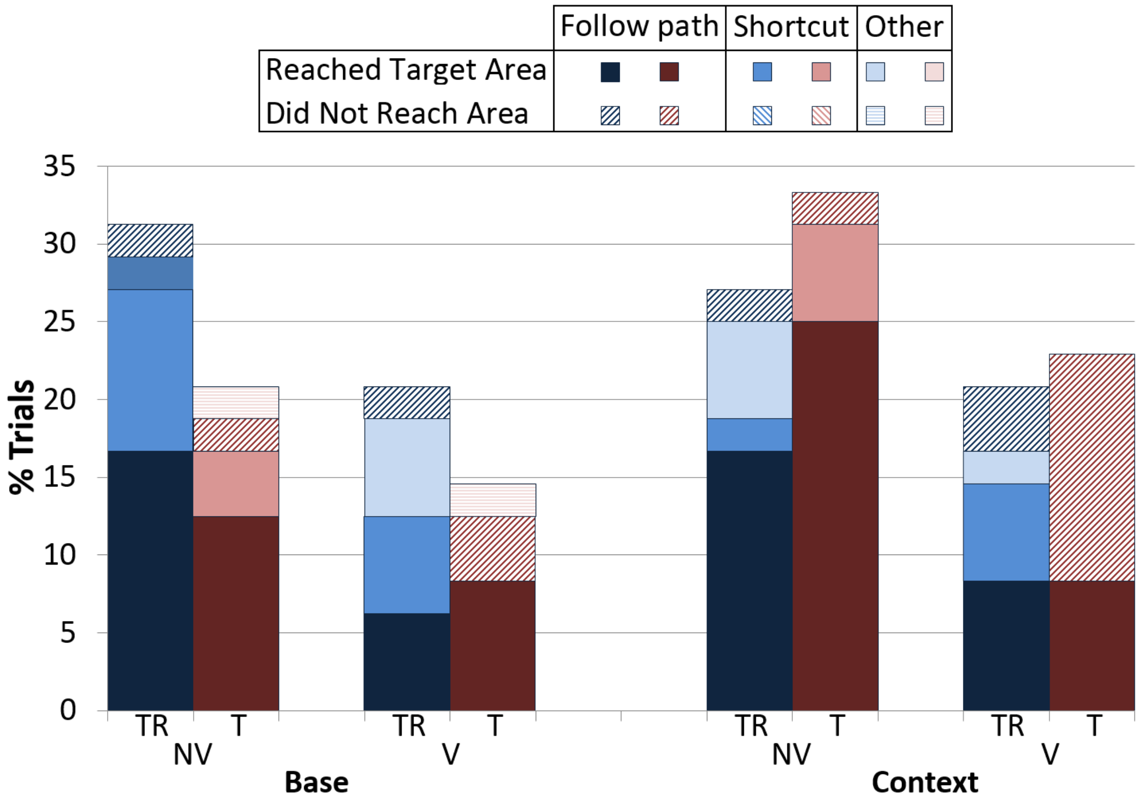

Figure 4.7: Navigation method in trials where participants failed to find the target for each combination of views, target visibility (visible V, not visible NV) and movement type (translate & rotate TR, translate T).

some translate paths but the normally direct routes these took to the target would mask this.

There was also a three way interaction between navigation method, visibility of target and movement type, $F(2, 9) = 5.98$, $p < .025$. This was due to participants attempting to follow the master's path with approximately the same frequency for both translate and translate & rotate paths when the target was not originally visible on the wall-sized display as well as for translate & rotate paths when the target was visible on the wall-sized display (*M* between 67 and 81%). However, participants almost exclusively attempted to follow translate paths when the target was visible on the wall-sized display (*M* = 98%) with almost no instances of shortcut or other strategies.

An increase in participants' following the original path seemed to have a positive influence on the success for translate movements in the base view condition whilst trials in the context view condition seemed less dependent on participants being able to follow the path. For example, when the target was visible, translate & rotate trials were more successful than translate and more participants took a short-cut than followed the path.

Trials where the participants failed to complete the trial were investigated to understand the severity of, and reasons behind the failure. Participants behaviour was classified in the same way as the successful trials with the addition of whether they managed to get to the approximate area of the target (within 3000m) or not. ANOVAs were not performed due to the small number of participants (2 out of the 12) that had failures in all combination of conditions.

Figure 4.7 shows that participants generally managed to navigate to the approximate area of the target, despite committing an error before getting there. In trials where participants did not reach the target area, the most frequent behaviour was that they attempted to follow the original path but committed an error from which they could not recover. Trials including a short-cut very rarely missed the approximate target area.

4.2.3 Discussion

Contrary to H1.1, a wide FOV context view did not significantly affect participants' performance through the metrics of success, time taken or distance travelled. This may have been because any advantage in terms of knowing where one travelled that was gained by the contextual display (showing the master's FOV) was negated by having to attend to a third view.

H1.2 was also largely shown to be false, with no overall difference in failure rate between translate and translate & rotate paths, even though the former just required participants to travel in straight lines. A statistically significant difference was found for the initial visibility of the target location, but the magnitude of this difference was small.

Despite this the experiment did result in evidence of non-trivial amounts of breakdown between wall-sized and desktop display users, with participants failing to find the target in 24% of trials. The method of navigation to the target was similar for successful and failure trials, apart from more errors being made in attempting to follow the original path in failure trials. These results indicate that participants often reached a target's region, but then had difficulty finding it. However, the extent to which participants were certain that they were in the correct region remains unknown.

Fatigue could be one factor that is masking some of the differences that were expected in this study. Although rest was provided between each of the blocks of trials, the repetitive nature of the task may have led to participants concentrating less in later blocks. This is apparent from figure 4.8 which shows participants improving from the training set in the first condition they experienced, but deteriorating in the second condition. This fatigue also confirms that the nature of navigating in CVEs is mentally taxing even without additional data analysis tasks, with performance beginning to suffer even over the course of an hour and a half.

The similarity of some of the areas within the data as well as the high amount of optic flow provided by the textured surfaces are factors that could have influenced the ability of users to navigate in this environment. However, this was a worthwhile trade off in providing a scenario which participants readily accepted and a dataset that they had some degree of familiarity with how to navigate. The textures used,

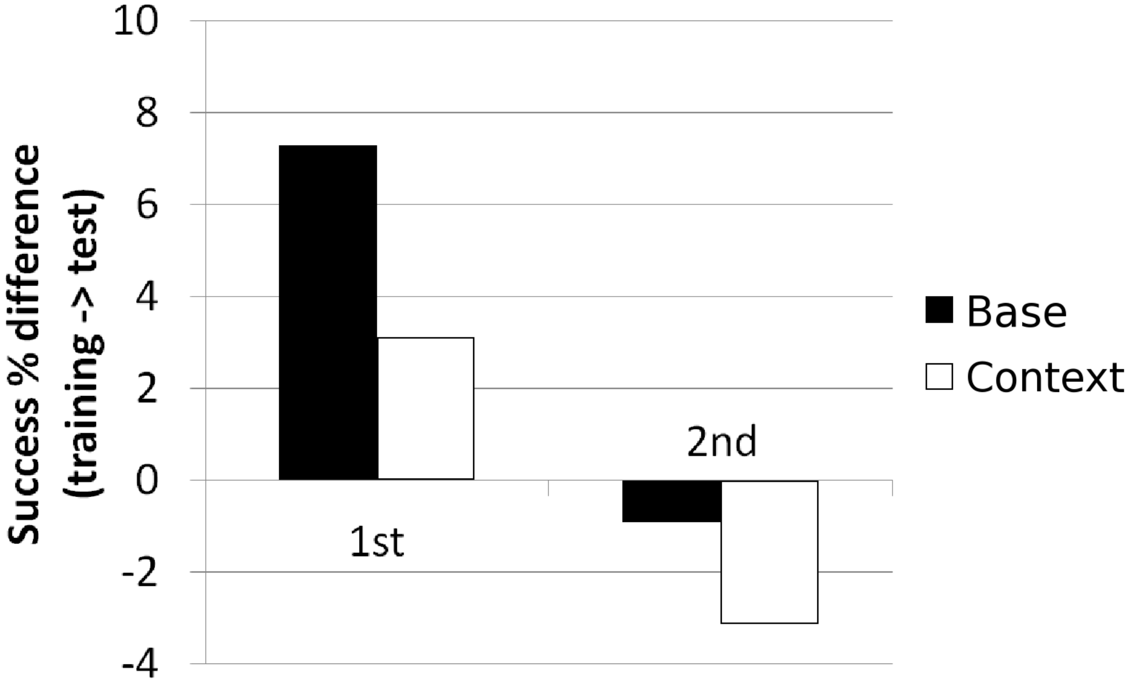

Figure 4.8: Difference in success rate between training and test blocks when a view condition was used 1st or 2nd.

although providing contrasting colours throughout, were of a reasonably low resolution when viewed closely which did create difficulties for some users when travelling close to the ground. The topological features (ridges) could have encouraged participants to fly low to see these clearly rather than at a higher altitude and therefore contextual manner.

Despite these issues it is notable that so many failures occurred even though the paths all comprised of flowing movements that lasted a few minutes at most. In a real collaborative application, navigation is likely to take place over an extended period of time, with the master sometimes making sudden changes to where they were travelling. This is likely to increase common ground breakdown, compared with the already non-trivial level reported in the present study.

4.3 Experiment 2: Free form navigation to a number of targets

The results of the previous experiment indicated that expanding the scope of the participants' task would give greater insight into elements of their behaviour which can be addressed to improve navigational performance. A second experiment was devised to investigate instances where the breakdowns within this scenario happen more frequently and either stop a user from being able to complete their task, or require them to adopt significantly different behaviours.

These issues were investigated by extending the task performed in the previous experiment so that participants had to navigate towards multiple targets per path and the master's navigation was controlled by a gamepad or mouse. Increasing the number of targets within the path allowed the path complexity to be increased whilst ensuring that the findings could be related and contrasted with those from the previous experiment. The mouse was introduced for the master's paths for two reasons: firstly, to further investigate the effect of the wall-sized display user using features that are outside of the desktop display view for navigation by allowing them to discretely travel to these points, and secondly to validate to what degree these findings are valid across different interaction devices.

An experiment with a within-participants design was conducted to investigate the frequency and severity of navigation errors that occur during master-slave navigation utilizing a wall-sized and a desktop display. Some participants acted as master users (henceforth termed users) using a wall-sized display whose paths were pre-recorded. Each user recorded half of their paths while using a gamepad to navigate and half while using a mouse. Both interfaces are known to be effective for navigating wall-sized displays [9, 122].

The remaining participants acted as slaves (henceforth termed participants) and used a desktop display to observe where the master 'took' them (from a start point to either one or three targets), and then navigate themselves from the start to the target(s). Our hypotheses were as follows:

H2.1 The discrete nature of mouse movements will cause greater common ground breakdown than the continuous nature of the gamepad movement.

H2.2 Increasing the demands of the task (i.e., the number of targets) will increase common ground breakdown, because of the additional cognitive load [8].

H2.3 Interruptions and inconsistencies such as finding targets in a different order to the master will increase the common ground breakdown.

4.3.1 Method

4.3.1.1 Factors

The within-participants' factors for this experiment were:

- Number of targets to find (one vs. three)
- Type of interaction device used by the master (mouse vs. gamepad)

The number of targets was varied between one and three to investigate how extended navigation and a more demanding task would affect the errors that occur, testing H2.2. Mouse and gamepad interaction allow investigation of displacement and velocity-based controls respectively, testing H2.1.

4.3.1.2 Measures

The following measures were taken during each trial:

- Success/ failure in finding each target
- Time elapsed during the trial
- Time elapsed before finding each target
- Order in which targets were found compared to the master's path
- Position and orientation per frame

The first two measures allowed performance to be assessed, the next two measures allowed hypothesis H2.3 to be tested while the last measure allowed participants' navigational behaviour to be investigated. Instances of common ground breakdown were observable through reduced navigational performance and/ or differences in navigational behaviour.

4.3.1.3 Participants

Twelve slave participants (three females and nine males, mean age = 22.8 years, *SD* = 3.1) and four master users (one female and three males, mean age = 24.0 years, *SD* = 2.2) took part in the experiment. The order in which participants viewed gamepad vs. mouse movement was counterbalanced and the ordering of one and three target trials was randomised.

4.3.1.4 Materials

The dataset, wall-sized displays used by the master and desktop display used by the slave participants were the same as in the first experiment. Gamepad navigation was also the same as in the previous experiment and was the sole form of control provided to slave participants. Interaction via the mouse was enabled in half of the master's paths and is described in further detail in section 3.6.

4.3.1.5 Procedure - Master

The movements of master users from start points to the required targets using the wall-sized display were recorded, intending to simulate a situation where a user is familiar with a dataset and is interested in particular features within it. The user's navigation is therefore expected to be smooth and direct as opposed to during data exploration where navigation is more varied. Pre-recorded paths rather than live navigation ensured that each slave participant was subjected to exactly the same movements.

Recording was conducted in one session for each user, which was split into three stages; a training stage, recording for one interface (e.g., mouse) followed by recording the second interface (e.g., gamepad). Each of these stages was separated by a short break and contained 16 trials in each (split evenly between one and three target trials). Movements controlled through the gamepad resulted in more continuous movement than those produced via mouse control where the user remains stationary for larger periods of time (see figure 4.9).

Each trial began by allowing familiarisation with the targets on a given path by highlighting each so they were visible from the start position. Users navigated

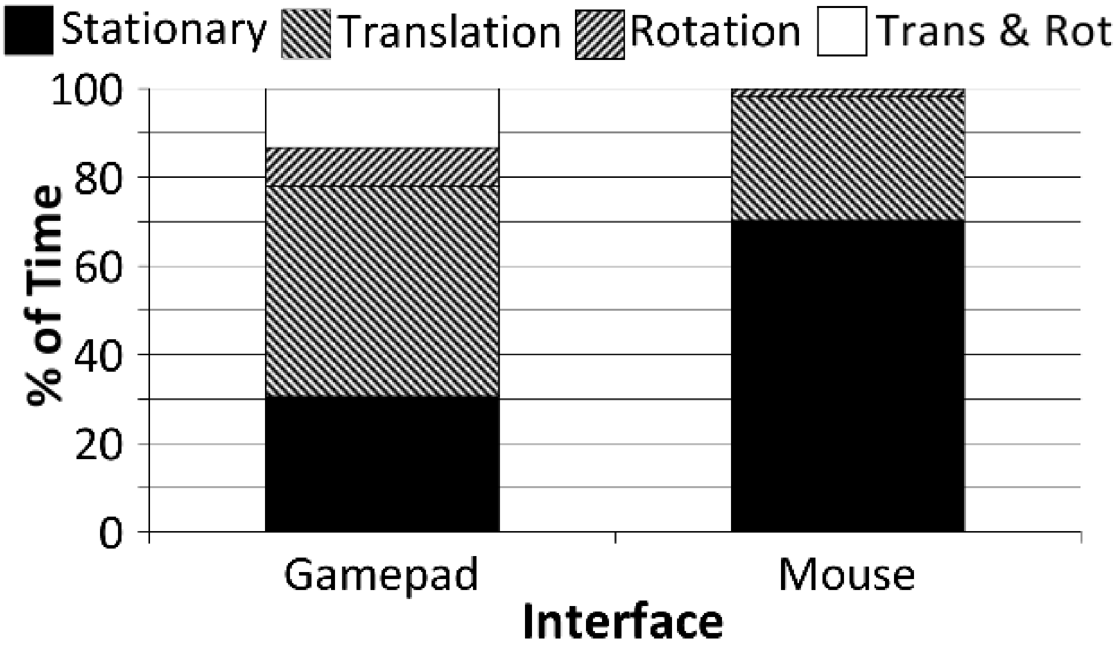

Figure 4.9: Proportion of time performing certain movements by the wall-sized display master users during the 32 test trials.

between these targets, typically for less than three minutes, and when ready the recording stage began through a button press on the gamepad, which returned the user to the original position and removed the highlighting of the targets. They were here asked to find the targets (now as shown in figure 4.2) before confirming each through pressing a button on the gamepad. The target was required to be within the outlined visible area of the desktop display when they confirmed sighting of each target. Throughout both the familiarisation and recording stage of each trial the targets were clearly highlighted on the overview map.

Paths that were not direct (e.g., the same target was found twice or took an excessive amount of time (>two minutes)) were excluded. The two sets of paths shown to the slave participants were formed by pseudo-randomly selecting one of the remaining paths from each user in turn so each contributed eight gamepad and eight mouse paths per set. The test trials' characteristics are shown in table 4.2.

Interface	1 Target Mean (SD)	3 Targets Mean (SD)
Gamepad	10880 metres (4316) 11.3 seconds (5.6)	23335 metres (8212) 29.6 seconds (8.6)
Mouse	11095 metres (6252) 12.0 seconds (11.7)	25890 metres (6252) 47 seconds (21.1)

Table 4.2: Average distance and duration of the 32 test trials.

4.3.1.6 Procedure - Slave

The slave participants' procedure was divided into four stages. First, each participant practised navigating around the landscape by spending up to five minutes navigating between six clearly-visible targets. The next stage involved a set of 16 training trials which were presented in a random order, with eight of the trials involving gamepad master movement and eight involving mouse master movement. In each eight, four were trials with one target and four were trials with three targets. Five minutes rest was given before the test trials (stages three and four), which had identical formats. Slave participants viewed one master interface's movements (e.g., gamepad) in one stage, and the other master interface (e.g., mouse) in the latter. Each stage contained 16 trials, split equally between one target and three targets trials, which were presented in a random order and separated by a five minute rest between the two stages.

The procedure for each trial was similar to the first experiment except participants had to confirm sight of each target (by pressing a button within 5000 metres of the target) and the cut off time was extended to three minutes (double the time of the longest trial). Participants were instructed during training to check the target was not a decoy before confirming each target.

4.3.2 Results

This section first analyses the general performance of participants at the task, before identifying the major factors that indicate the occurrence of navigational error. Following this, the severity of navigational error within failed trials is analysed.

4.3.2.1 General Performance

The percentage of successful trials was analysed using a repeated measures analysis of variance (ANOVA) that treated the number of targets (one vs. three) and master's interface (mouse vs. gamepad) as within-participants factors. Participants were significantly more successful at one target trials (*M* = 90%, *SD* = 11.5) than when three targets had to be found (*M* = 55%, *SD* = 22.3) ($F(1, 11) = 46.3$, $p < .01$). There was no significant difference in success between the mouse and gamepad trials (*M* = 69% vs. 76% $F(1, 11) = 2.54$, $p > .05$). Although success was affected by the number of targets in a trial, participants were similarly likely to complete a one-target trial as to find at least one target in the three-target trials (*M* = 96%, *SD* = 7.1). Success in finding subsequent targets decreased steadily (*M* = 80%, *SD* = 18.4 for the second target, *M* = 56%, *SD* = 22.7 for the third).

These results do not support our first hypothesis, as there was little difference due to the master interface, but do support our second hypothesis that extended and more demanding tasks are more likely to result in navigational error. The remainder of the results only consider three-target trials.

4.3.2.2 Factors Influencing Success

Our third hypothesis predicted participants would be less successful if their navigation was interrupted or if they did not find the targets in the same order as the master. Approximately 80% of the targets were found within 40 seconds of the start of a trial or after finding a previous target, irrespective of whether a participant was looking for the 1st, 2nd or 3rd target, or the master's interface used in the trial. Interruptions were therefore analysed by splitting the trials into two sets; one set comprised trials where every target was either found in less than 40 seconds or not found at all, and the other set comprised trials where at least one target took 40s or longer to find.

These sets were analysed using a repeated measures ANOVA that treated the time threshold ($<$40s vs. $\geq$40s) and master's interface (mouse vs. gamepad) as within-participants factors. Only significant effects adjusted for bonferroni corrections ($p < .013$) are reported for this and the following ANOVA. Participants were significantly less successful on interrupted trials ($F(1, 11) = 11.87$, $p < .01$), but

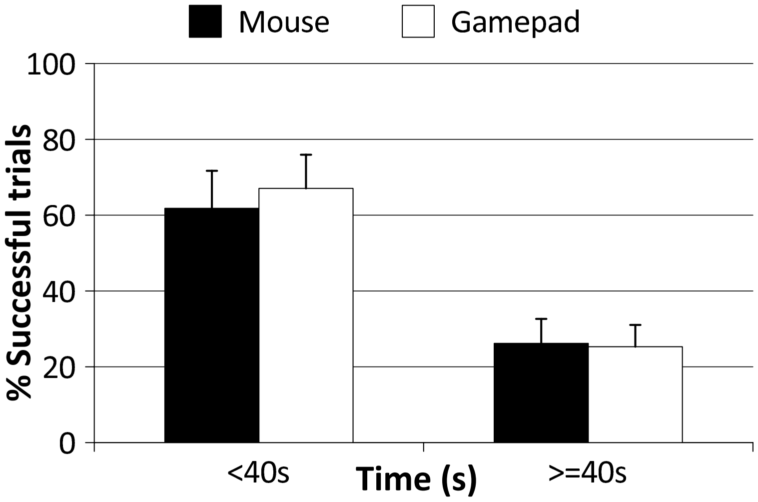

Figure 4.10: Success rate in three target trials when no preceding targets took more than 40 seconds to find vs. when any preceding target took more than 40 seconds to find. Error bars show standard error of the mean.

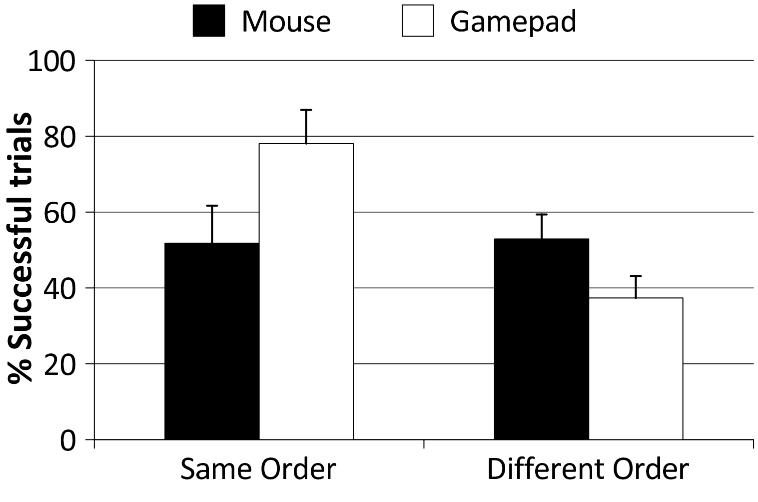

Figure 4.11: Success rate in three target trials when all preceding targets were found in the same order to the master vs. when any preceding target was found in a different order to the master. Error bars show standard error of the mean.

there was no significant effect of master interface (*F*(1, 11) = 0.85, $p > .025$) (see figure 4.10).

To investigate the effect of finding targets in a different order, trials were divided into two sets according to whether or not a participant found all the targets in the same order as the master, up to the point of failure. A repeated measures ANOVA which treated the order of targets (same vs. different to the master) and master interface (mouse vs. gamepad) as within participant factors showed that participants were significantly less successful in trials where targets were found in a different order than when found in the same order as the master (*F*(1, 11) = 10.06, $p < .01$). However, there was no effect of master interface (*F*(1, 11) = 2.45, $p > .025$) (see figure 4.11). These results support our third hypothesis.

4.3.2.3 Classifying the Failures

To categorise the severity of failures, trials in which a participant did not find all three targets were classified according to whether participants got close to a remaining target in the 40 seconds after the last successfully found target. Participants were classified as getting close if they got within 3000 metres in the ground plane of the target, the same distance as used for this classification in the previous experiment (see section 4.2.2.2). In 59% of failures (26% of total trials) a participant got close to a remaining target, whilst failing to do so in the remaining 41% of failures (18% of total trials).

4.3.3 Discussion

The results from this experiment show H2.1 to be false as there was no significant difference in participants' success rate between mouse and gamepad trials. Firstly, this ensures that findings within this scenario can be applied across different interaction devices. Secondly, in combination with the results from Experiment 1 showing H1.2 to be false, these results provide evidence that the findings could be applied independently of the master navigation characteristics. These findings also provide further evidence to the unresolved issues of the effect of passive master/ slave navigation [19, 90, 118, 133].

H2.2 was supported, with this experiment showing that initially trivial navigation (≈90% of trials found one target), became significantly more difficult as the complexity of the task increased (failures occurred in ≈45% of 3 target trials). In addition, H2.3 was also supported, with participants exhibiting difficulties in recovering from mistakes in navigation, such as interruptions or finding targets in a different order to the master, which further impacted their ability to complete the task.

Participants could complete the task either through route knowledge (recreating the master's route) or survey knowledge (travelling for a required distance at the required orientation) but both were frequently insufficient to reach the approximate area of a target (≈20% of trials). The frequency of these errors would distract one from their main analytical tasks in the scenario described in section 2.6 to such a level that this is currently not feasible.

Failures close to the target indicate moderate navigational error and may have been caused by participants not pausing to look around if they did not realize they were in the vicinity of a target, or missing targets that were within their view due to not being able to attend to the whole display [76]. Failures further away from the target indicate the occurrence of more severe navigational error, despite a global overview map always being provided. The risk of missing landmarks from the primary view while viewing the map, coupled with the difficulty in integrating knowledge gained from the two views together [102], may have dissuaded participants from doing this.

4.4 Summary

This chapter presents two experiments which classify the frequency, severity and conditions of common ground breakdown within the scenario presented in section 2.6. Experiment 1 investigated what effect (if any) a wide FOV context view (H1.1) and variances in navigational task complexity (H1.2) had on participants' performance. Both of these hypotheses were shown to be false, partially due to the task's simplicity providing little opportunity for variation in navigational behaviour between participants.

Therefore, Experiment 2 investigated the effect of master interface differences (H2.1), increased task complexity (H2.2) and interruptions to the task (H2.3) on task performance. H2.1 was shown to be false, whilst both H2.2 and H2.3 were supported, overall showing a consistently high level of breakdown.

To summarise, the major characteristics of common ground breakdown in wall-sized to desktop collaboration were identified as follows:

- There is a non-trivial amount of common ground breakdown, which rapidly and steadily increases as users have to conduct more complex tasks resulting in participants failing to complete 45% of trials with multiple targets. This breakdown is across different types of master movement and input device.
- Participants managed to navigate towards the approximate area of the targets in approximately 60% of trials ending in failure even with the simplest of navigation aids. Breakdowns occurred more frequently when participants were trying to find the precise location of the target to complete the trial.
- Interruptions to common ground reduced task success.

The frequency of both moderate and severe navigational errors shows that participants need assistance in acquiring and utilising both route and survey knowledge (either of which could have been used to complete this task) to reduce the quantity and/or severity of the errors that occur. There have been a number of aids developed that aid navigation by providing landmarks, trails of previous routes or additional views of the environment. The role of landmarks in navigation is generally accepted [107]. Trails can be used to address errors in integrating route and survey knowledge where participants repeatedly revisit the same areas of an environment and aid long term route knowledge retention [98]. Providing additional or alternative views, with a wide FOV and a local overview map, can aid route [34, 102] and survey knowledge respectively. The following chapters will explore the effect that these aids have on the errors demonstrated in this chapter.

Chapter 5

Additional Views

5.1 Introduction

This chapter will investigate the effect the navigation aid of additional views (suggested in section 4.4) will have.

5.2 Development

As discussed in section 2.4.1, views can range across a continuum from human's-eye (purely egocentric) to a north-up global map (purely allocentric). The two views currently provided of the environment are examples at the extreme ends of this spectrum. This leads to users' difficulties in being able to combine knowledge from the two together into a comprehensive cognitive map. If views that are in a more moderate location within this spectrum are provided they should allow this assimilation of knowledge to be performed more quickly and accurately.

5.2.1 Context (Egocentric) view

An egocentric view can provide a greater sense of contextual position through increasing the FOV of the environment that it displays [105]. A FOV that is identical to that of the wall-sized display used in Experiments 1 and 2 would allow a participant to view landscape features (see section 2.6) that would otherwise only be visible to

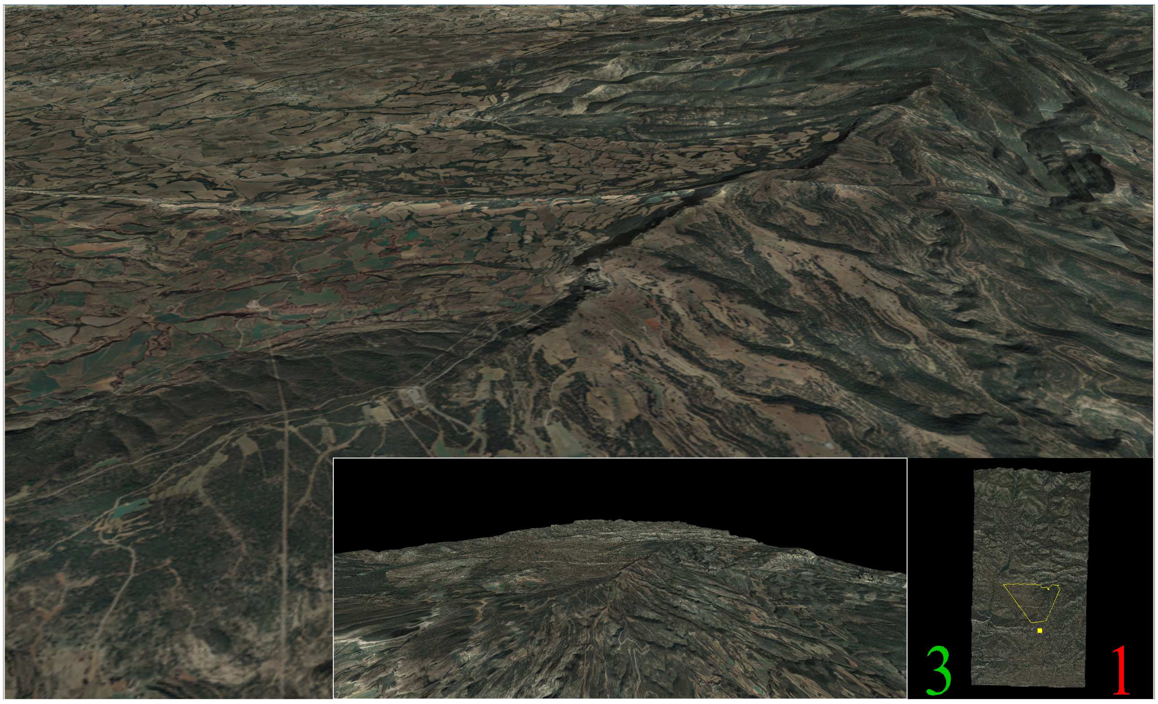

Figure 5.1: Example of the display layout with the context view. The context view is outlined in white along the bottom of the display alongside the global map. The two numbers overlaid on the global map show the participant's current trial (right number) and how many targets they still have to find in the current trial (left number).

the master user, providing greater fidelity to a participant's route knowledge [121]. However, the desktop user would not be able to view these at the same resolution as on the wall-sized display. This may be equally detrimental to the common ground as it would be impossible to see detailed elements that the wall-sized display user is currently analysing.

Therefore to provide the user the opportunity to see both detail and context within the view the standard and large FOV views can be combined in two ways: through distortion or as separate views. We have chosen to use separate views to avoid the difficulties associated with the judging distances and object shape as they cross the border of the distortion. An example of the implemented view is shown in figure 5.1.

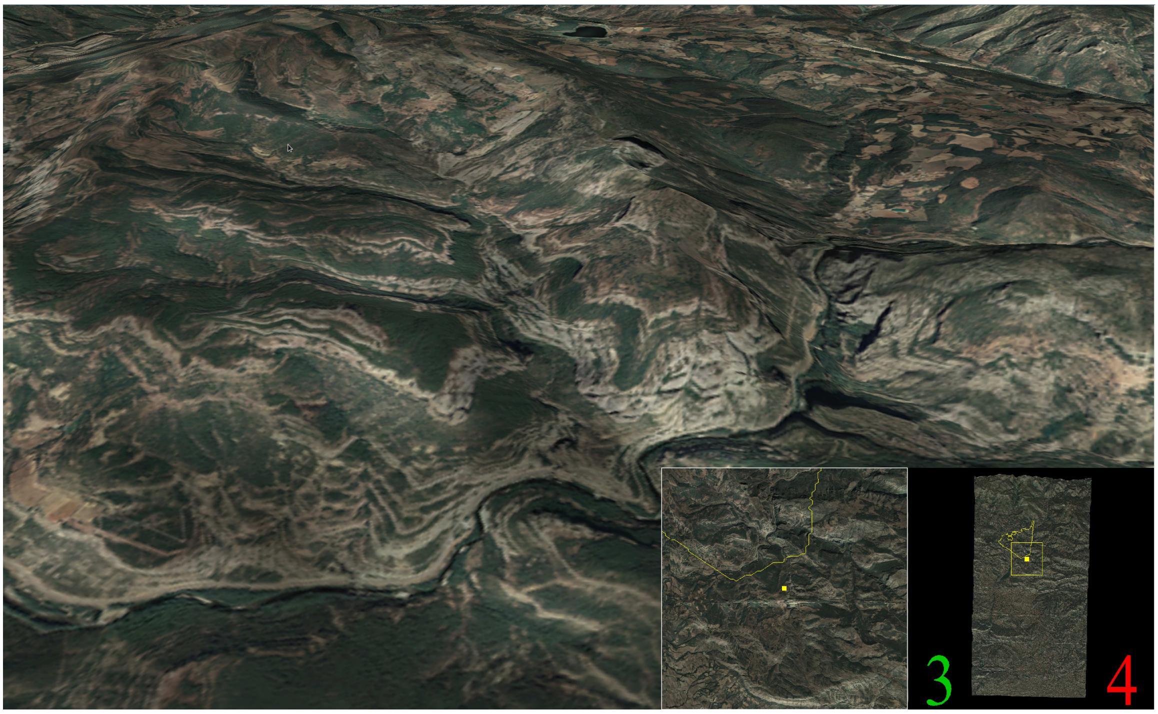

Figure 5.2: Example of the display layout with the local map view. The local view is outlined in white on the bottom of the display alongside the global map. The two numbers overlaid on the global map show the participant's current trial (right number) and how many targets they still have to find in the current trial (left number).

5.2.2 Map (Allocentric) view

Differences in orientation and scale between the main and overview map could cause users difficulties in combining spatial knowledge across the two [102, 121]. Depending on the orientation of the view provided the user may have to perform extensive mental rotations to combine knowledge from the two views. This could be helped by providing a map which maintains the same orientation as the user's main view, but may only be useful in a limited number of circumstances than the you are here information provided by the current global map.

The differences in scale between the main view and the overview map is a more frequent issue and can be addressed through the use of a local map at a scale between the two. A suitable scale would allow the user to see the landscape features used for navigation (see section 2.6) that would be distinguishable if orientated towards them, even when they are not visible in the main view. In addition, a map at this scale would provide more information to mentally rotate the features that are visible within their view to the map orientation rather than only relying on mentally rotating large landmarks. These benefits should result in participant's exhibiting greater survey knowledge [121]. An example of such a local map is shown in figure 5.2.

5.3 Experiment 3: Navigation when utilising additional views

This experiment investigated the effect context and local overview map views have on the frequency and severity of navigational errors that occur during master-slave navigation. The context view provided the participants with the same view area as the wall-sized display master but at a reduced scale. The local map had a fixed north up orientation, was centred on the user and scaled so that objects up to 3000m away were visible. The two views (as shown in figures 5.1 and 5.2) both supplemented a base view that was the same as that provided in the last experiment (see figure 4.2). We had the following hypotheses regarding this experiment:

H3.1 The context view should improve participants' route knowledge (see section 5.2.1) and result in a higher success rate at finding targets.

H3.2 The local view should improve participants' survey knowledge (see section 5.2.2) and result in a higher success rate at finding targets.

H3.3 Participants should be more succesful when using the same interface as the master (gamepad) than when using an interface that prevents them from replicating the masters movements (e.g., mouse).

5.3.1 Method

5.3.1.1 Factors

Only trials with three targets were used and resulted in two within-participant factors:

- The type of view used by the slave (base vs. context vs. local)
- The master interface (gamepad vs. mouse)

The first factor allows H3.1 and H3.2 to be tested while the second factor allows H3.3 to be tested.

5.3.1.2 Measures

The following measures were taken during each trial:

- Success/ failure in finding each target
- Time elapsed during the trial
- Time elapsed before finding each target
- Order in which targets were found compared to the master's path
- Position and orientation per frame

The first two measures allowed performance to be assessed, the next two measures allowed interruptions and inconsistencies to following the master path to be analysed while the last measure allowed participants' navigational behaviour to be investigated. Instances of common ground breakdown were observable through reduced navigational performance and/ or differences in navigational behaviour.

5.3.1.3 Participants

Twelve slave participants (four females and eight males, mean age = 24.6 years, *SD* = 5.5) and four master users (one female and three males, mean age = 24.0 years, *SD* = 2.2) took part in the experiment. The slave view was viewed in a Latin Square arrangement whilst the master interface was presented in a random order within a set.

Interface	Set A Mean (SD)	Set B Mean (SD)	Set C Mean (SD)
Gamepad	29526 metres (9817) 37.4 seconds (15.9)	32156 metres (9250) 32.1 seconds (9.4)	34059 metres (11800) 29.7 seconds (13.9)
Mouse	34898 metres (15499) 49.0 seconds (19.6)	32983 metres (8859) 49.0 seconds (21.0)	36772 metres (13239) 59.4 seconds (31.99)

Table 5.1: Average distance and duration of the 48 test trials.

5.3.1.4 Materials

The materials for this experiment consisted of the same geophysics dataset, wall-sized master system used to record the paths and desktop slave system used by participants as described in section 4.3.1.4.

5.3.1.5 Procedure - Master

Master recording was conducted in two sessions, each of which was split into two stages. A training stage allowed users to become familiar with the process for each trial using both interfaces. After a short break the next stage recorded paths with one device (e.g., gamepad) then the other (e.g., mouse). The first session consisted of 16 (eight per interface) training trials and 16 (eight per interface) recorded trials, the second session eight (four per interface) training trials and 32 (16 per interface) recorded trials. The characteristics of the movement produced by gamepad and mouse control were similar to those described in section 4.3.1.5.

Each trial was conducted in the same manner as Experiment 2. Paths that were not direct (e.g., the same target was found twice or took an excessive amount of time (>two minutes)) were excluded. The three sets of paths shown to the slave participants were formed by pseudo-randomly selecting one of the remaining paths from each user in turn so each contributed eight gamepad and eight mouse paths per set. The test trials' characteristics are shown in table 5.1.

5.3.1.6 Procedure - Slave

The slave participants' procedure was divided into three sessions, each of which consisted of three stages. The procedure for each session was the same except for the view interface used (base, local or context view). First each participant practised using the interface by spending up to ten minutes navigating between eight clearly-visible targets two times. The next stage involved a set of eight training trials presented in a random order, with four gamepad trials and four mouse trials. Five minutes rest was given before stage three, which included the test trials. This stage contained 16 trials, split equally between gamepad and mouse trials, which were presented in a random order.

The procedure for each trial was the same as Experiment 2. In each trial, the time and movements that participants took were recorded for subsequent analysis.

5.3.2 Results

This section first reports the performance of participants under the different conditions to determine what differences (if any) are significant before explaining these findings through analysis of participants' behaviour.

5.3.2.1 Performance

The views provided were expected to result in changes to the quantity of failures that participants exhibited, so the percentage of successful trials was analysed using a repeated measures ANOVA that treated the master interface (mouse vs. gamepad) and slave view (base vs. local vs. context) as within-participant factors. This analysis showed that there were no significant effects of master interface ($F(1, 11) = 0.02$, $p > .05$) or slave view ($F(2, 10) = 0.45$, $p > .05$) on success. Means varied from 54% ($SD = 26.3$) for the mouse master interface and context view to 59% ($SD = 25.0$) for the mouse master interface and base view. Success in finding additional targets decreased between the first ($M = 94\%$, $SD = 6.9$), second ($M = 81\%$, $SD = 14.8$) and third target ($M = 57\%$, $SD = 20.7$) in a similar manner to the second experiment (see section 4.3.2).

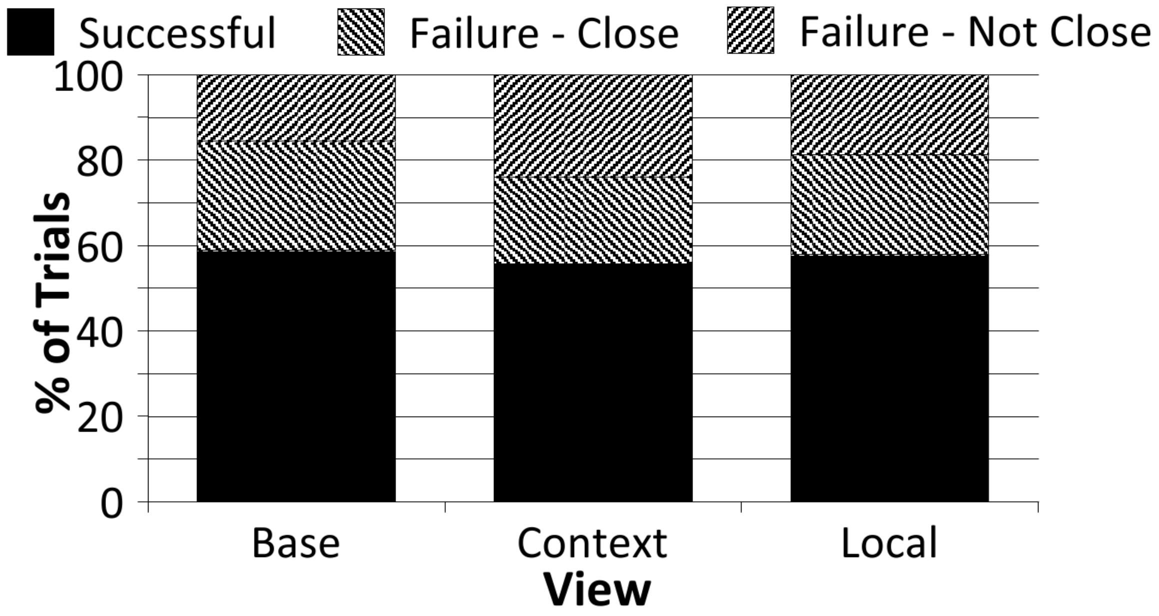

Figure 5.3: Percentage of trials for each view condition in three categories. In the first category participants were successful, in the second they failed but got close to a target they needed to find and in the third they failed as well as not getting close to a target they needed to find.

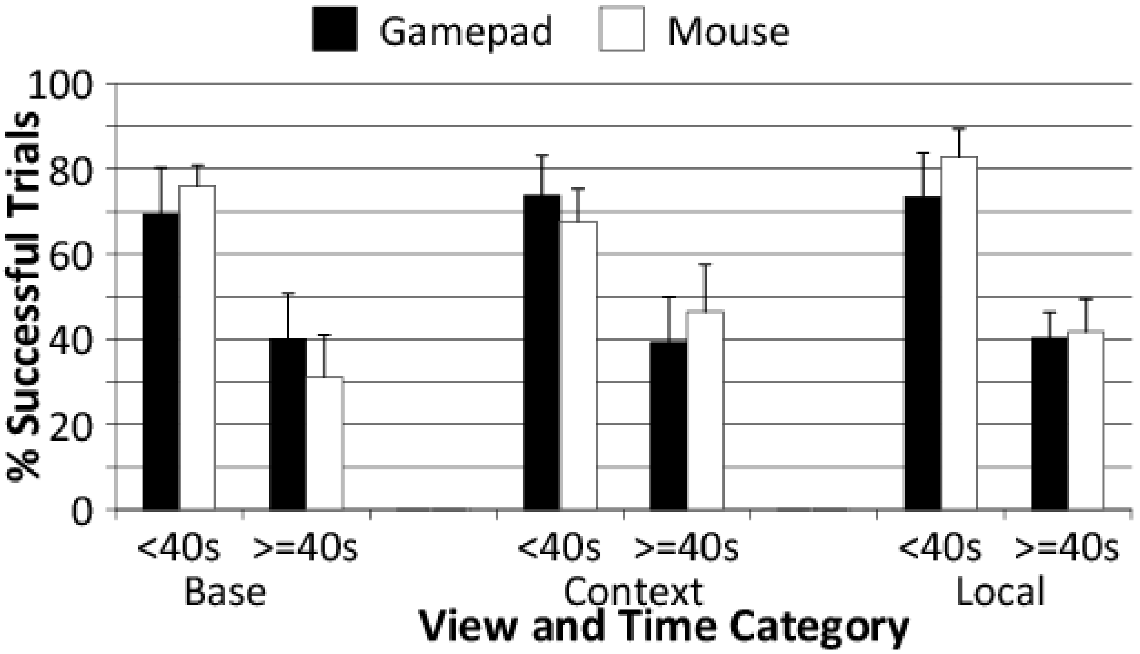

Figure 5.4: Success rate when no preceding targets took 40 seconds or more to find vs. when any preceding target took 40 seconds or more to find across the view conditions. Error bars show standard error of the mean.

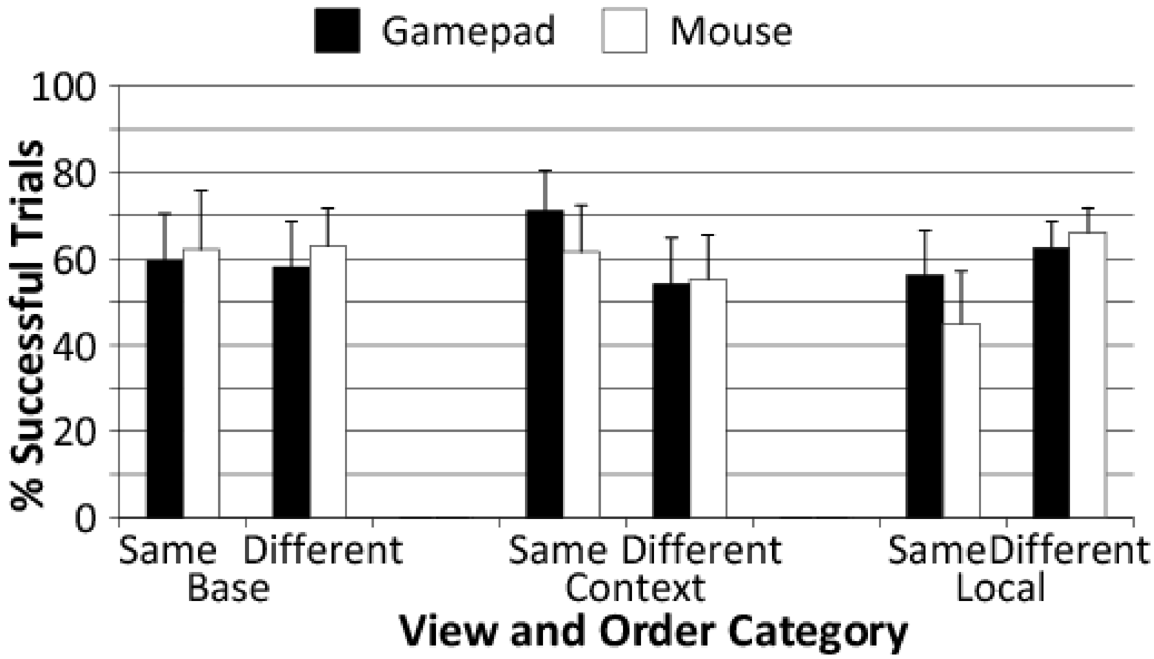

Figure 5.5: Success rate when preceding targets were found in the same order as the master vs. when any preceding target was found in a different order to the master across the view conditions. Error bars show standard error of the mean.

The types of failure that occurred were classified in the same way as the previous experiment (see section 4.3.2.3) and the percentage of trials with these failures was analysed using a repeated measures ANOVA that treated the type of failure (close vs. not close to the target) and slave view (base vs. local vs. context) as within-participant factors. There was no significant effect on the type of failure ($F(1, 11) = 0.58$, $p > .05$) or view ($F(2, 10) = 0.30$, $p > .05$) (see figure 5.3).

In the previous experiment targets that were found following an interruption ($>$40s looking for a target) or in a different order to the master significantly decreased the likelihood of successfully completing the trial. The effect of interruptions was tested by splitting the trials into two sets as described in section 4.3.2.2. A repeated measures ANOVA that treated the type of view (base vs. context vs. local), master interface (gamepad vs. mouse) and amount of time to find the target ($<$40s vs. $\geq$40s) as within-participant factors and showed that there was a significant effect on success when an interruption occurred ($F(1, 8) = 71.98$, $p < .01$) similarly to the first experiment (see figure 5.4). However, there were no significant effects of view ($F(2, 7) = 0.50$, $p > .05$) or master interface ($F(1, 8) = 0.00$, $p > .05$).

The effect of the order in which the targets were found on success was also investigated through a repeated measures ANOVA that treated the type of view (base vs. context vs. local), master interface (gamepad vs. mouse) and order the targets were found (same vs. different to the master) as within participants factors. This showed that there was no significant effect of view ($F(2, 10) = 0.16$, $p > .05$), master interface ($F(1, 11) = 0.50$, $p > .05$) and unlike the first experiment no significant effect of target order ($F(1, 11) = 1.00$, $p > .05$) (see figure 5.5).

5.3.2.2 Behaviour

One of the few differences between the different view conditions was the increase in failures where participants did not get close to the target in the context view condition compared to the base view condition. After analysing participants' movements during these failures three main patterns of behaviour became apparent and are described below:

1. The participant set off towards the target correctly but did not reach it, stopping or turning away too early, indicating an error in the distance they thought they

had travelled [107]. This was deemed to occur when the vector from the last successfully found target to the point where the participant is 3000m away from the target was <45 degrees from any of the vectors from the last successfully found target to the remaining targets (see figure 5.6).

2. The participant accurately completed a route towards an invalid target, either a route to a target they have previously visited (see figure 5.7) or a route to go between two targets starting from the wrong target (see figure 5.8). This pattern shows participants utilised a route-based strategy to travel between targets but had insufficient survey knowledge to recognize that they were travelling towards the wrong area. This was deemed to occur when the participant visited (came within 3000m of) a previously found target or a point whose vector from the last successfully found target is that between any of the targets in the trials.

3. Some other error occurred that could not be classified.

These classifications were applied to the movements previously classified as failures where participants did not get close to the target in section 5.3.2.1 and instances of each are quantified in figure 5.9. The increase in these failures is shown to be mainly due to participants accurately completing a route to an incorrect target.

If the additional views are providing participants with a sufficient contextual view of the environment, the main view can be used for viewing the detail of the environment more frequently. In this situation the participant would be positioned nearer to the point of the terrain in the center of their view (the point of interest (POI) for this analysis).

This distance was measured across the first 40s of searching for each target and analysed using a repeated measures ANOVA which treated the type of view (base vs. context vs. local) and distance from the target (close (<3000m away) vs. approaching (≥3000m away)) as a within-participant factor. This analysis showed that participants were significantly closer to their POI during the context view condition than when the base or local view conditions ($F(2, 10) = 5.45$, $p < .05$), and significantly closer to their POI when close to the target than when approaching it ($F(1,$

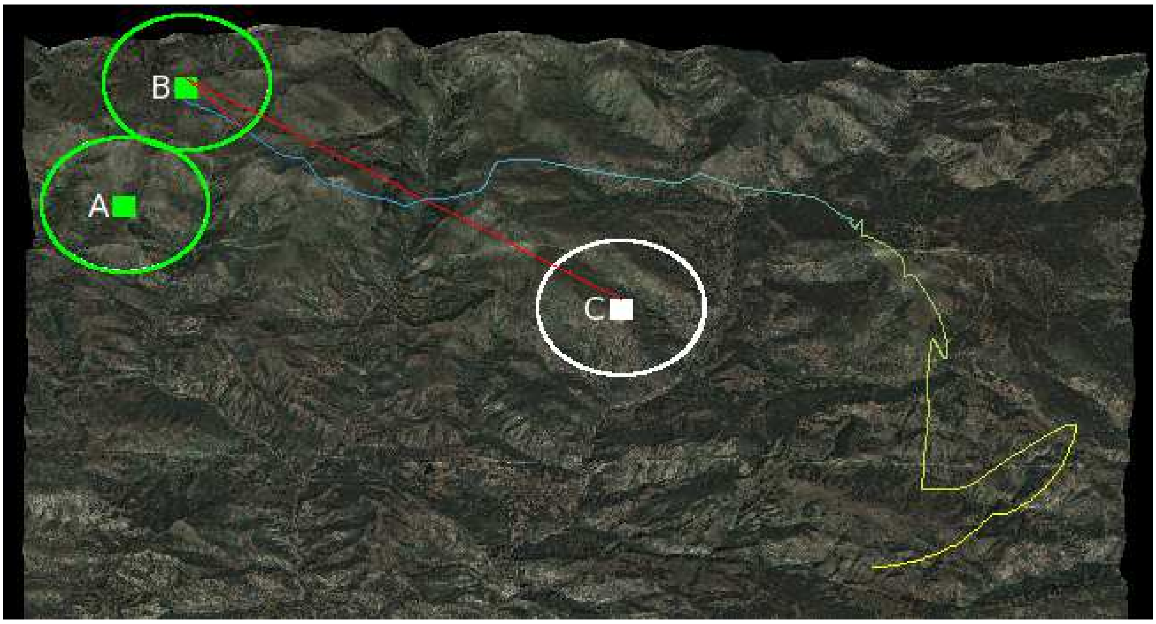

Figure 5.6: Examples of when a participant does not get close to a target that needs to be found after setting off correctly in the direction of the remaining target. The participant's path after the last successfully found target is shown as progressing from blue to yellow. The target they are trying to find is highlighted in white and previously found targets are highlighted in green. The circle around each target has a 3000m radius.

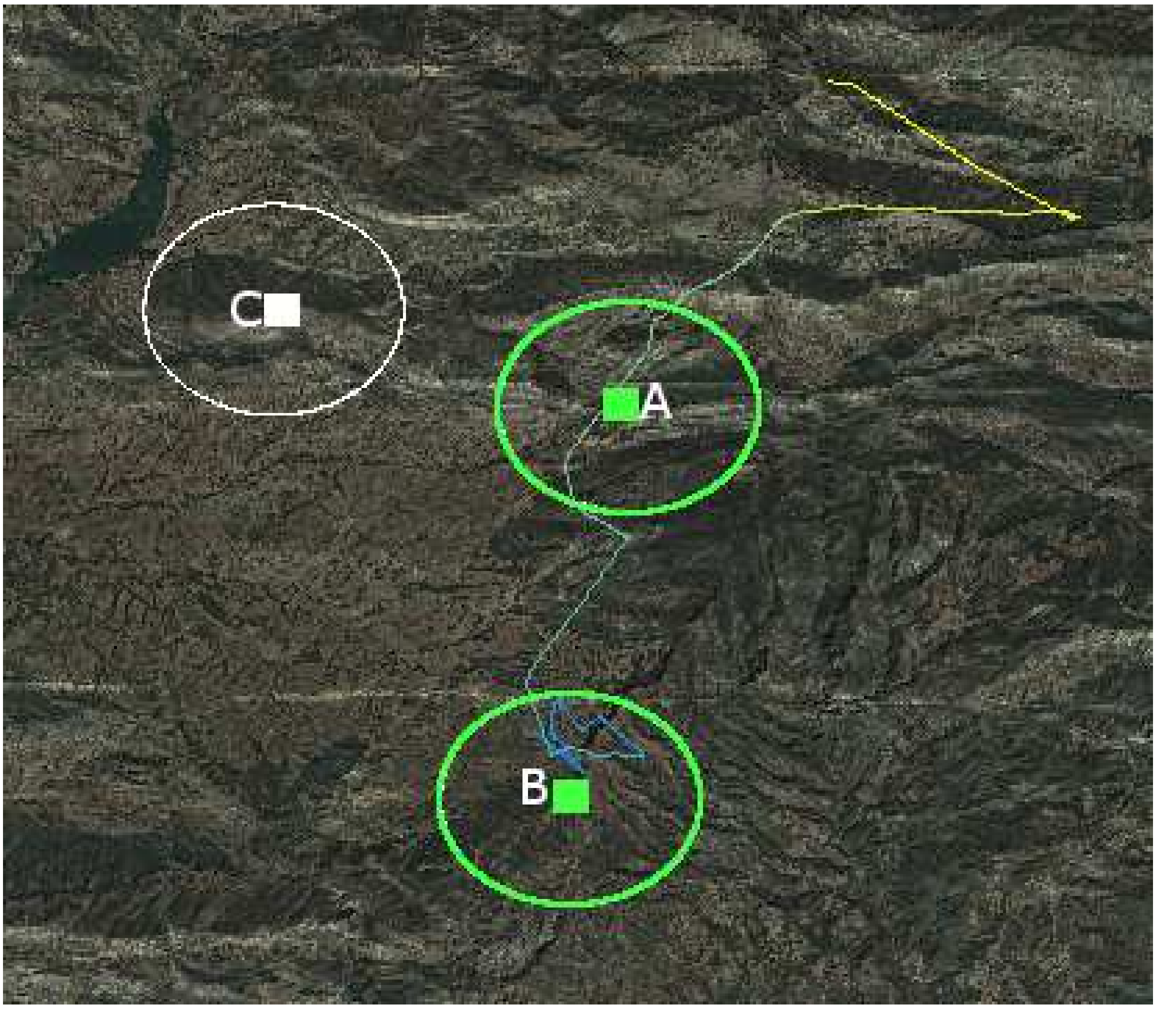

Figure 5.7: Examples of when a participant does not get close to a target that needs to be found because they travel towards a previously found target . The participant's path after the last successfully found target is shown as progressing from blue to yellow. The target they are trying to find is highlighted in white and previously found targets are highlighted in green. The circle around each target has a 3000m radius.

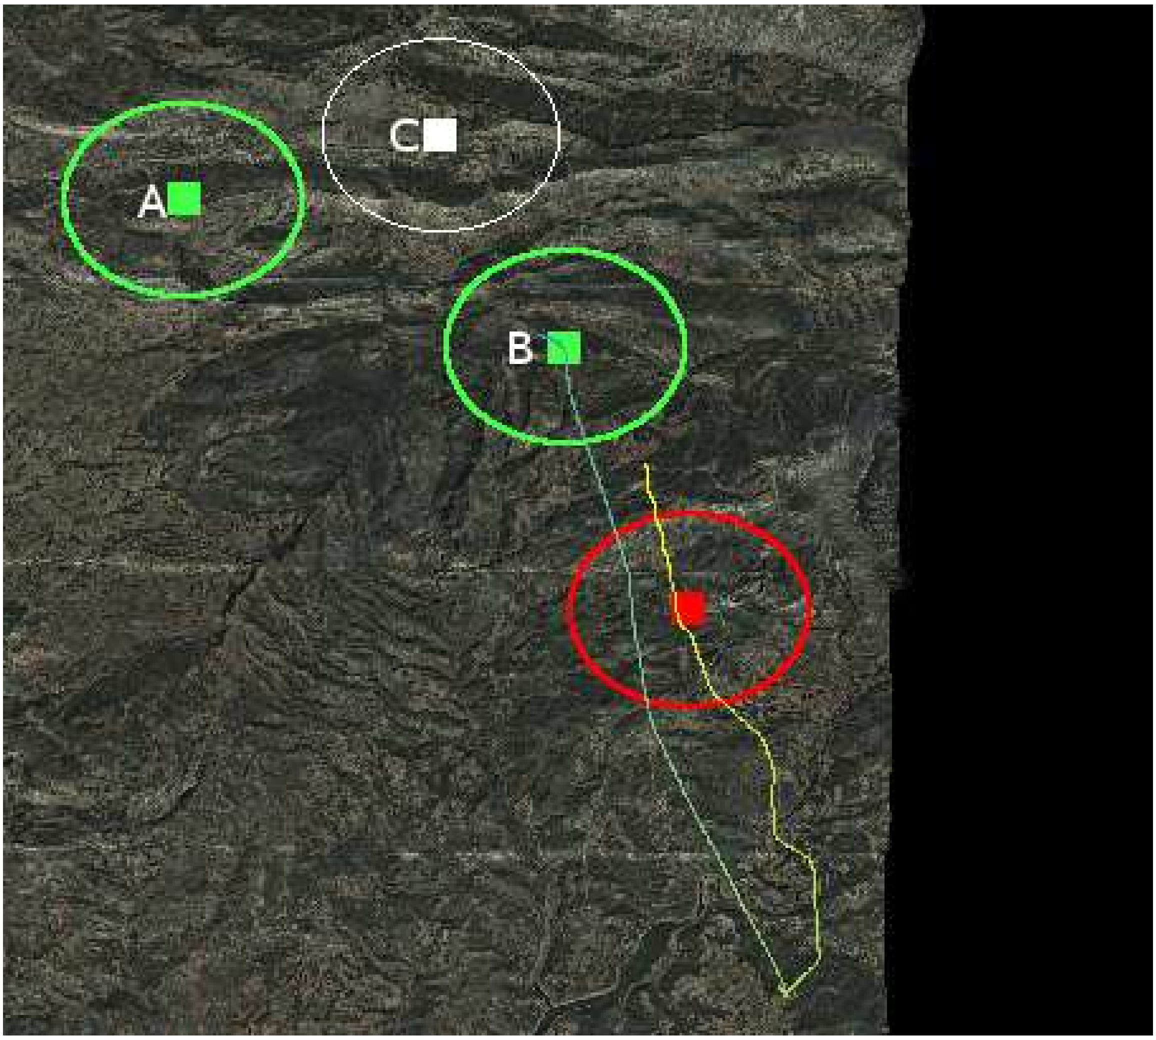

Figure 5.8: Examples of when a participant does not get close to a target that needs to be found because they travel towards a imagined target. The participant's path after the last successfully found target is shown as progressing from blue to yellow. The target they are trying to find is highlighted in white and previously found targets are highlighted in green. The circle around each target has a 3000m radius. The participant traverses a path as if they are going to target B from target C but travels from target B instead of C (The position where they would expect the target to be is highlighted in red).

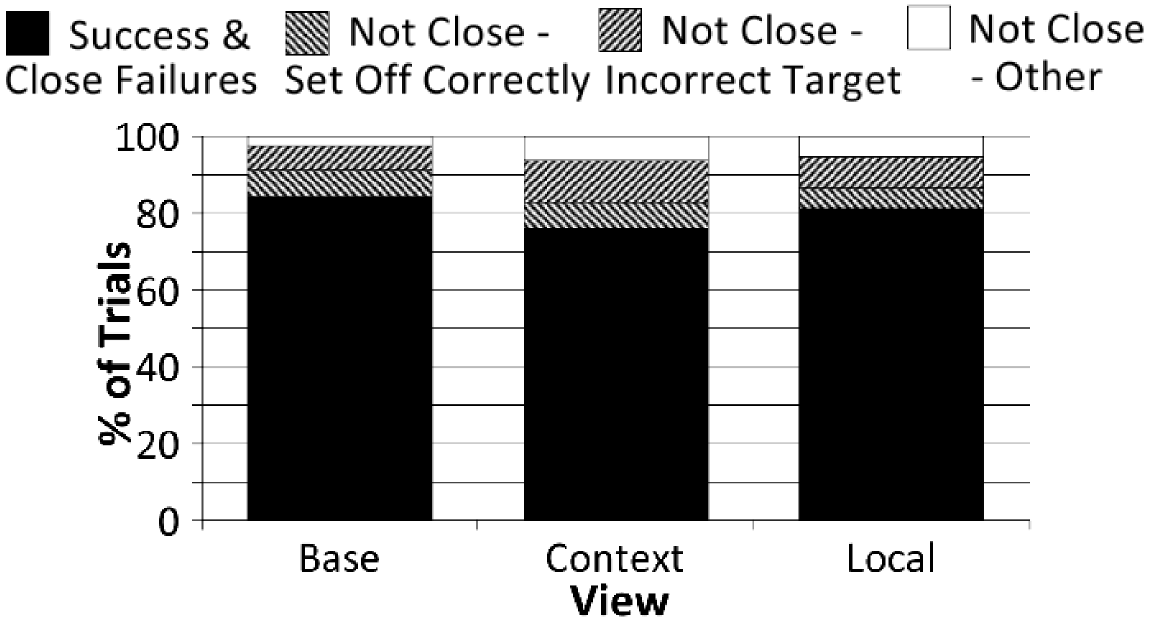

Figure 5.9: Percentage of trials for each view condition in four categories. In the first category participants were either successful or got close to a target they needed to find whilst failing the trial. In the remaining categories the participant did not get close to the target. In the second category participants set off correctly towards the target, in the third they visited an incorrect target and in the fourth category their movements could not be classified into the previous categories.

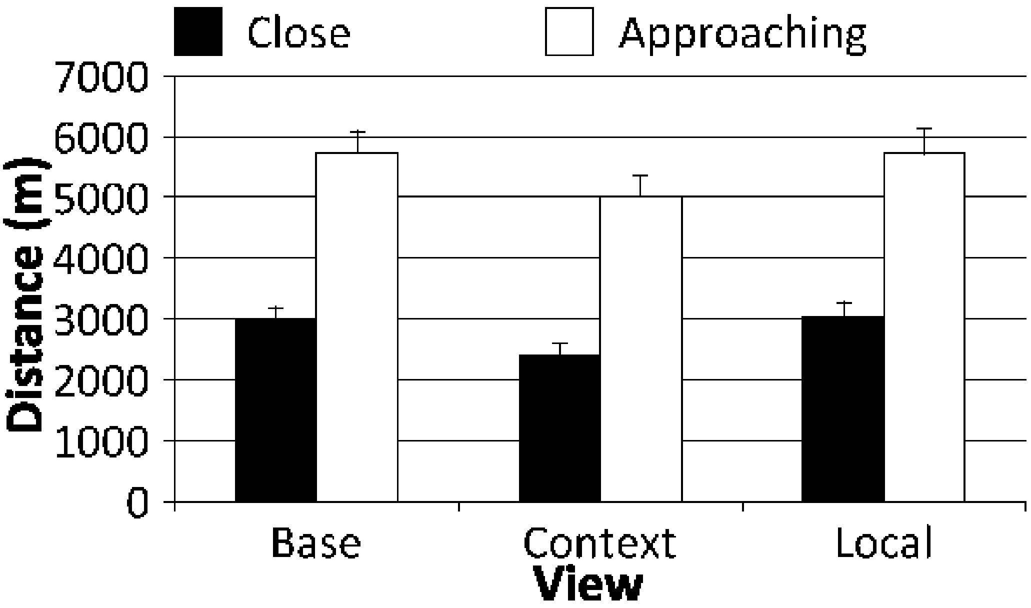

Figure 5.10: Average distance between the participants' position and the centre of the view on the terrain when approaching the target vs. when close to the target for each view condition. Error bars show standard error of the mean.

11) = 150.24, $p < .01$) (see figure 5.10). The difference caused by the view condition indicates that participants were using the context view to contextually view the environment as discussed above.

5.3.3 Discussion

It is surprising that context and local map views had no beneficial effect when these types of view have been effective or suggested in other scenarios [34, 69, 102], showing H3.1 and H3.2 to be false. The interface used by the master had no significant effects throughout the experiment, further showing H3.3 to be false.

There was evidence that participants used the context view to see the environment in both detail and context simultaneously (mimicking the use of a wall-size display), as well as to adopt a route based strategy more frequently, a strategy requiring more accurate wayfinding than a distance and orientation approach (see section 2.3.1). However, these differences ultimately resulted in more severe failures occurring, most of which saw participants taking a route towards a previous target or starting the route from the wrong area. This behaviour could be reduced through marking master and/or slave movements with trails to distinguish previously visited and unvisited areas [98].

The local view had little effect on both the frequency of navigational errors and the behaviour of participants. This view may have been ineffective or not used due to difficulties in recognizing nearby elements of the environment, a problem avoided in previous examples by clearly highlighting specific landmarks [102]. The fixed and limited range of the map may have also caused difficulties in relating landmarks within the local map to those found in either the main view or the global map.

5.4 Summary

This chapter presents an experiment that evaluates the effect of a wide FOV context view (H3.1), local overview map (H3.2) and different master interface devices (H3.3) on the common ground breakdowns within wall-sized master to desktop slave collaborative navigation. Results demonstrated that neither view reduced the level of common ground breakdown exhibited by desktop users who performed similarly to the base condition. Behavioural data showed that participants used the context view in a manner similar to that which was hypothesised but still demonstrated similar types of task failure. Finally, master interface differences also had no significant effect on participant success rates.

Chapter 6

Path Visualisation

6.1 Introduction

An alternative aid to the multiple views presented in chapter 5 are path visualisations. These techniques can help users identify which locations have been previously visited to avoid them searching in unnecessary locations (an ineffective behaviour exhibited in the previous experiments). Information on where multiple users have visited should also guide them towards areas of the environment that are shown to be interesting (the location of the targets in this case). They can also be used to ensure that users do not repeatedly revisit the same areas of the environment whilst exploring themselves.

This chapter describes the process of developing and evaluating a design for a path visualisation that is suitable for the experiment scenario used previously. The design was informed by a technical evaluation of the design space, before it was investigated with an experiment to validate the effectiveness of the final visualisation.

6.2 Initial Development

Suitable path visualisations were developed by identifying key aspects of the possible design space through analysis of previous work (see section 2.4.2), and considering these in relation to the specific data being used in this scenario. A technical evaluation was conducted to narrow this design space further and the findings used

Variable			Levels	
Data	Where		Position	View
	Number of People		1	4
	Sampling Interval	Time	per frame (0.03s) to 1s	
		Grid Size	100m to 1000m	

Table 6.1: Design Space Variables.

to inform the final design of the path visualisations. The final designs were then implemented in a user study (see section 6.4) to evaluate their effectiveness.

6.2.1 Design Space Variables

To create a suitable path visualisation for this scenario the different variables of the design space need to be set appropriately for the scenario in which they are being used. This section will consider the suitability and parameter value of the variables extracted from the literature in section 2.4.2 for the final aid.

The recorded paths of the master participants from the previous studies were used to provide representative data and establish the main characteristics across variations in the input variables. Using this data limits the *who* and *what* data (see table 2.1) to individuals and movement respectively.

The chosen data limits the evaluation to single a user performing the tasks independently of each others without having to consider the complications introduced by a group conducting the task collaboratively. However, the data of multiple users can still be combined to explore scalability issues when larger quantities of data are used. The data chosen also limits the *what* data to movement actions as these were the only actions that could be performed in the previous studies. This leaves the *where* data types and the sampling intervals to be investigated, with the input variables that will be considered within this evaluation shown in table 6.1.

The *where* data of the path can utilise the position at which a user is in 3D space or what they are looking at (their view). Their view location will be considered as the point in the centre of their view on the landscape. Each recorded path needs to be considered as separately occurring events in time, but within each trial the time at which each of the points occurred has been recorded. This allows the specific time, order of, or frequency of points within a specified time period to be considered.

The frequency and resolution at which the data used for the visualisation is sampled both need to be set appropriately to its purpose (see section 2.4.2.1). The frequency of sampling could be on demand or at discrete, regular times. On-demand sampling will not be considered as it was not recorded as part of the original paths, would require additional master interaction and is traditionally seen as restrictive to otherwise successful systems [39]. However, a suitable frequency at which the data is discretely sampled must be determined to trade off between removing inconsequential movements within the path, and gaps between data points resulting in uncertainty between points. The data was previously recorded every frame (0.03s) so the data will be considered at this sampling frequency as well as frequencies of every 0.1s, 0.3s, 0.5s, 0.8s and 1.0s.

The resolution can be changed through the size of the sampling regions for the data. All the position data will be considered without the height component and the remaining two dimensions can be considered at different sizes to create square or rectangular ”bins” for the data. This introduces a grid structure to the data with the X and Y dimensions being kept equal so that square areas are used and varied between 100m and 1000m, in 100m increments.

6.2.2 Path Metrics

The resulting visualisation needs to allow users to cover as little of the environment as possible whilst allowing the depicted path to be easily followed. If a path visualisation allows users to immediately dismiss a subsection of the environment as being not relevant, because the original trailblazer did not travel there, then the visualisation will help users to navigate efficiently. The lower the coverage of a path, the larger the amount of the environment that can be dismissed in this way. A suitable metric to measure this is the number of sampled grid squares that the path populates as a percentage of the total number of sample grid squares within the environment. The example in figure 6.1a shows that 9 regions were traversed by the path out of the 36 regions in this area resulting in a coverage of 25%.

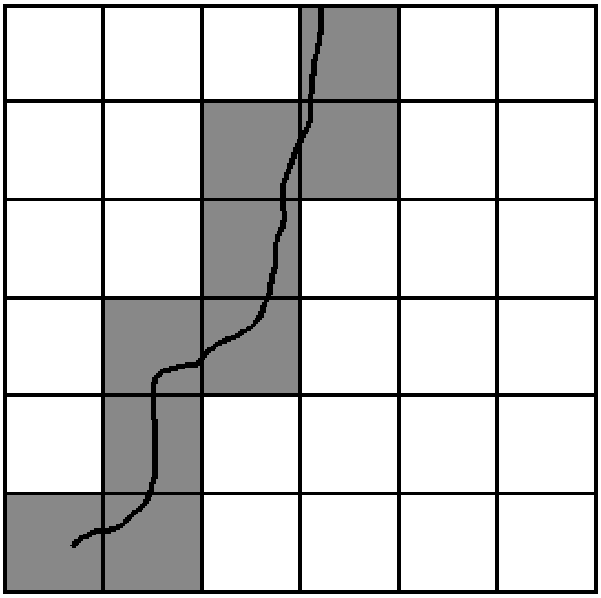

(a)

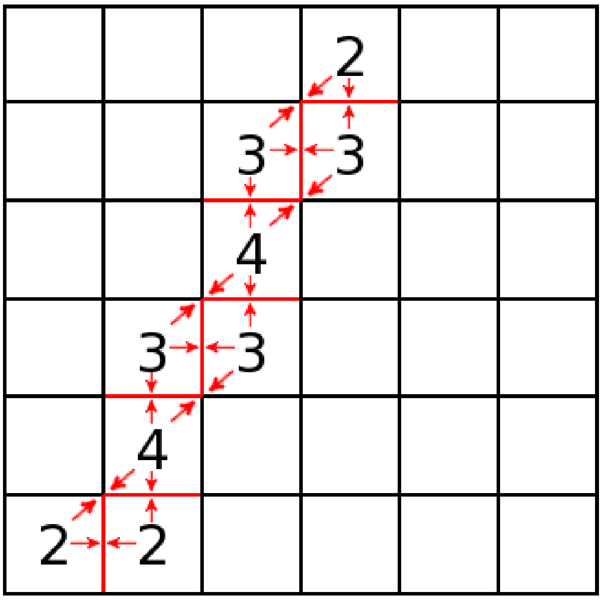

(b)

Figure 6.1: An example path that demonstrates the way in which coverage and adjacency metrics are calculated. The path and the sampled regions it passes through are shown in (a). The resultant adjacent regions that have been traversed are shown in (b) with a red arrow pointing between them and the total number of traversed adjacent regions for each region shown as a number in the centre of each region.

The manner in which a path is visually portrayed will affect how easily it can be followed. Ideally a path should be continuous but depending on the speed of movement and the sampling resolution of the original data there may be times when consecutive points are not in adjacent sampled regions. Equally, the adjacency of a square should be as near to 2 (one traversal into and one out of the region) as possible so there are few branches or potential crossings to complicate the path (see section 2.4.2). The metric for this is the average number of adjacent grid squares (above, above-right, right, below-right, below, below, left and above-left) for each of the grid squares that are covered by the path. In the example shown in figure 6.1b there are a total of 26 adjacent populated regions over 9 regions for an average populated adjacency of 2.89.

6.2.3 Technical Evaluation Method

The paths between targets from the second recording session of the four wall-sized display participants in Experiment 3 (see section 5.3) were used within this technical evaluation, resulting in two sets of 16 paths (32 in total) per participant, 128 paths between the four participants. Metrics were calculated (as described in section 6.2.2) for each set of targets initially using the path of a single participant (see figure 6.2a) to produce the single person results, and subsequently with the paths of all four participants (see figure 6.2b) for the multiple people results.

6.2.4 Technical Evaluation Results

The results are split into two sections, the first reports the results of a single participant's paths and the second the results from multiple participants' paths. Within each section, each factor within the design space is discussed in turn before considering any interactions.

6.2.4.1 Single Person

The coverage and the adjacency of the grid squares for a single user's path under the different design factors is shown in figure 6.3, with examples of the effect these factors have on an example path shown in figures 6.4 and 6.5. Sampling the data

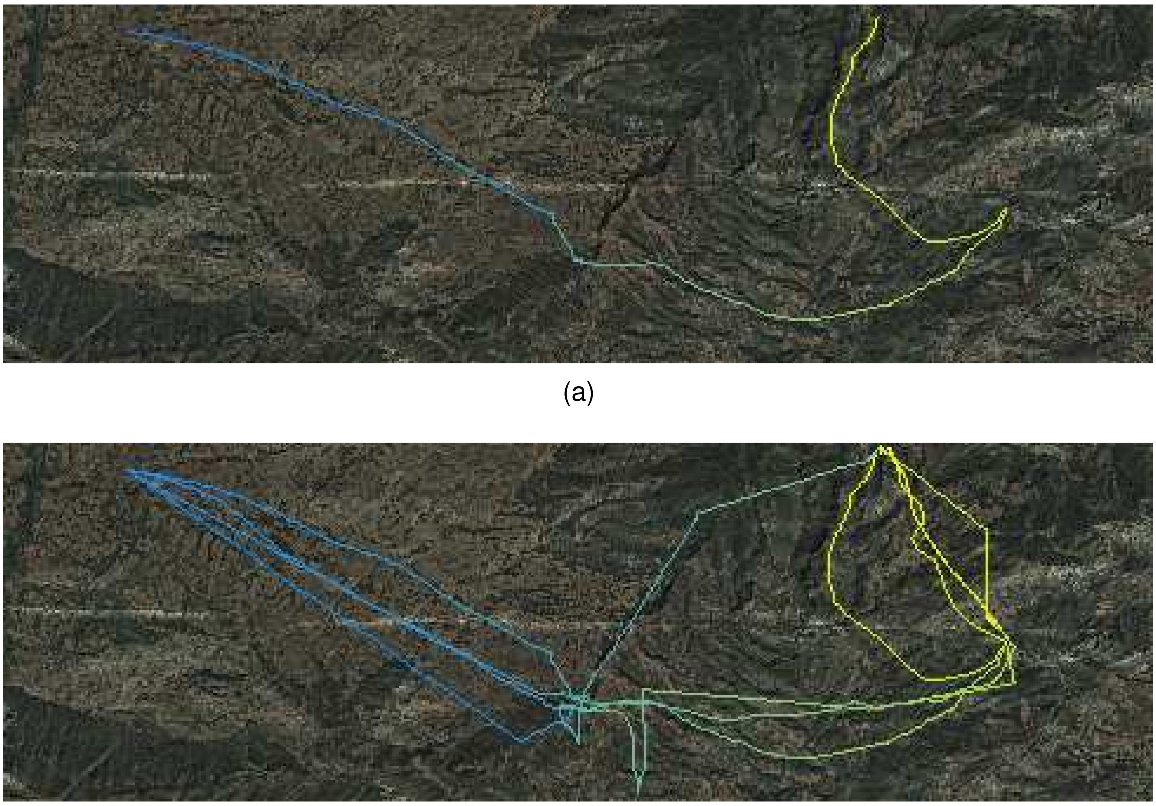

(a)

(b)

Figure 6.2: Examples of the paths used in the technical evaluation. A single wall-sized participant's path between three targets is shown in (a), with the path of four wall-sized participants between the same targets shown in (b). The paths are shown as a line going from blue to yellow as time progresses.

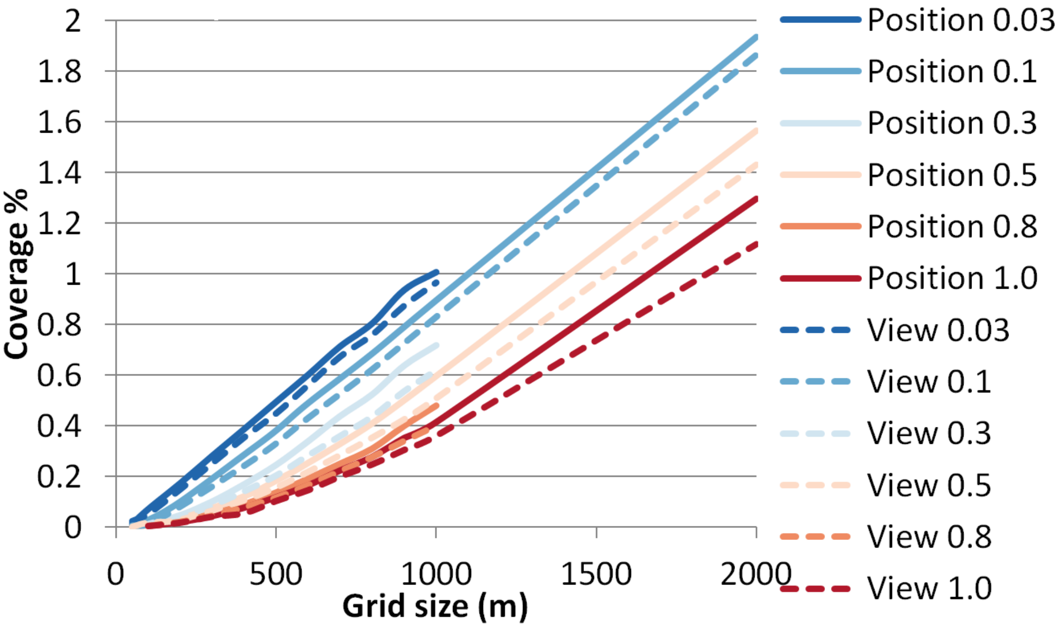

(a)

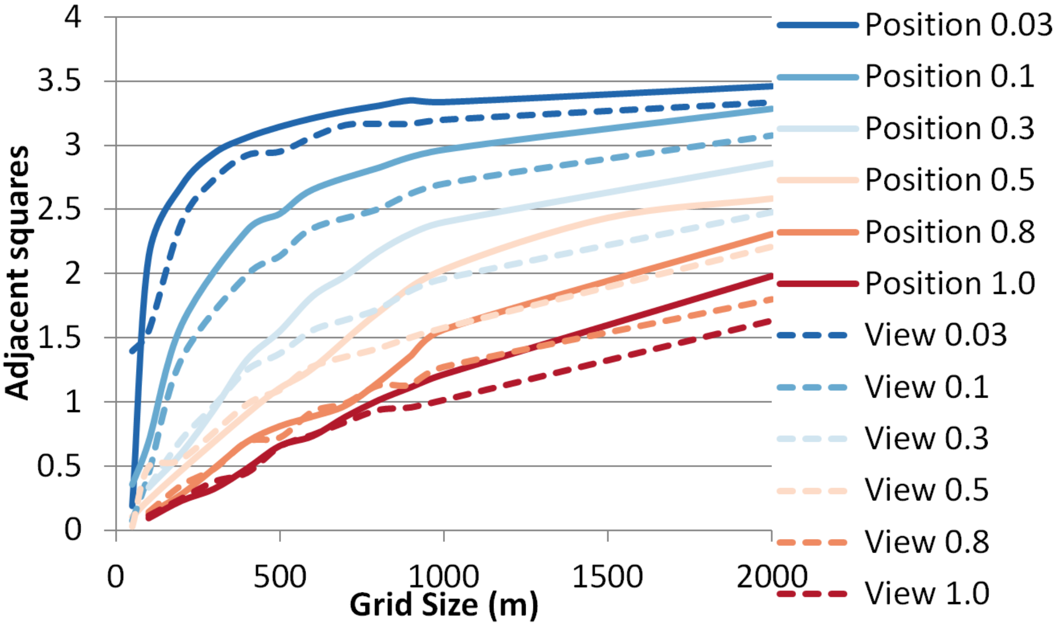

(b)

Figure 6.3: Results of analysing single person paths at different sampling intervals of time (between every frame (0.03s) and 1.0s) and grid size. The percentage of the environment covered by the path over a range of grid sizes is shown in (a). The average number of adjacent squares per populated square is shown in (b).

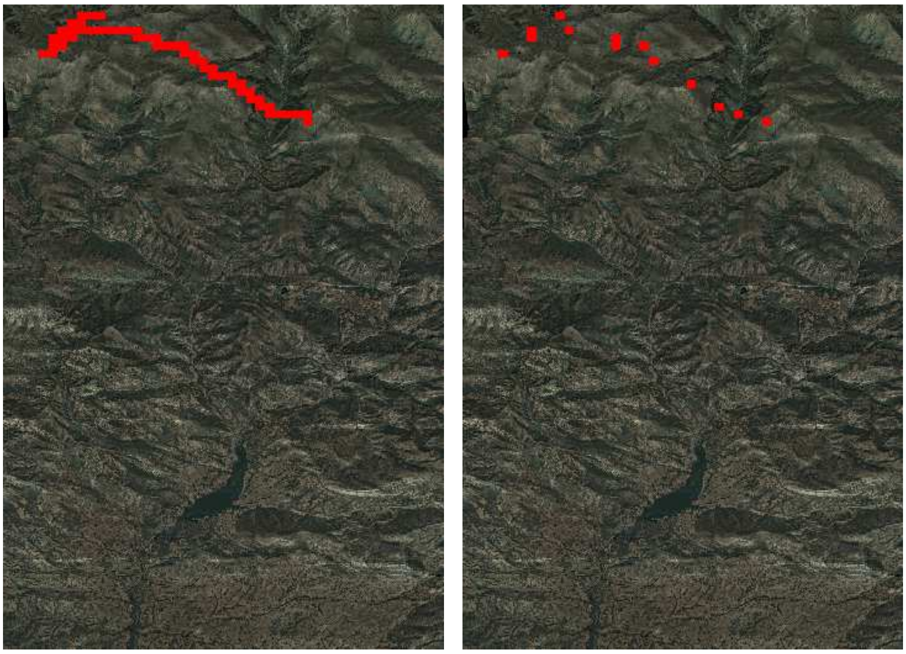

(a) Sampled every 0.03s at a grid size of 500m.

(b) Sampled every 1.0s at a grid size of 500m.

(c) Sampled every 0.03s at a grid size of 2000m.

(d) Sampled every 1.0s at a grid size of 2000m.

Figure 6.4: Examples of the position paths from a single person at different sampling frequencies and grid sizes. Each sub-figure has the same scale.

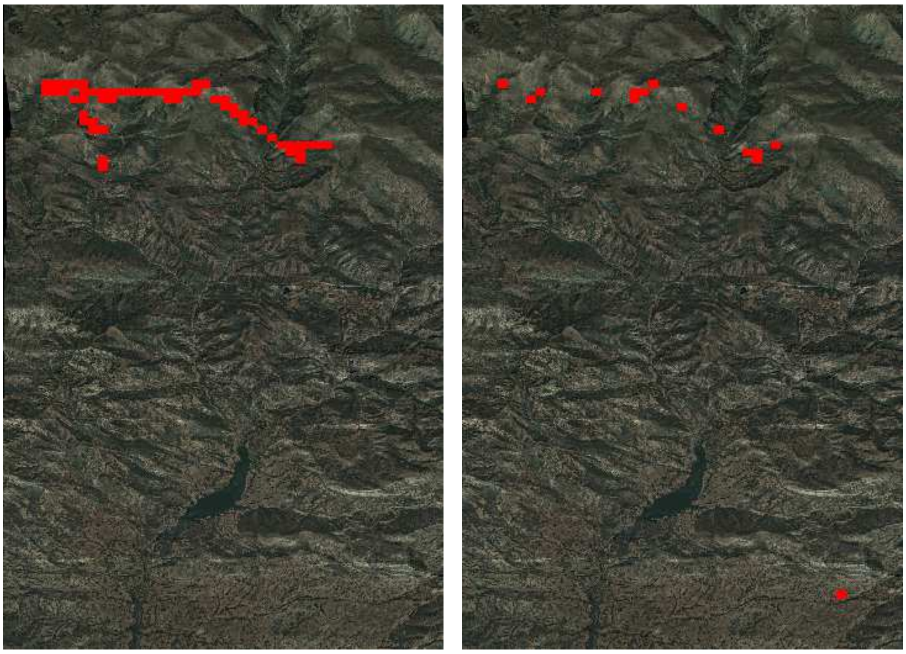

(a) Sampled every 0.03s at a grid size of 500m.

(b) Sampled every 1.0s at a grid size of 500m.

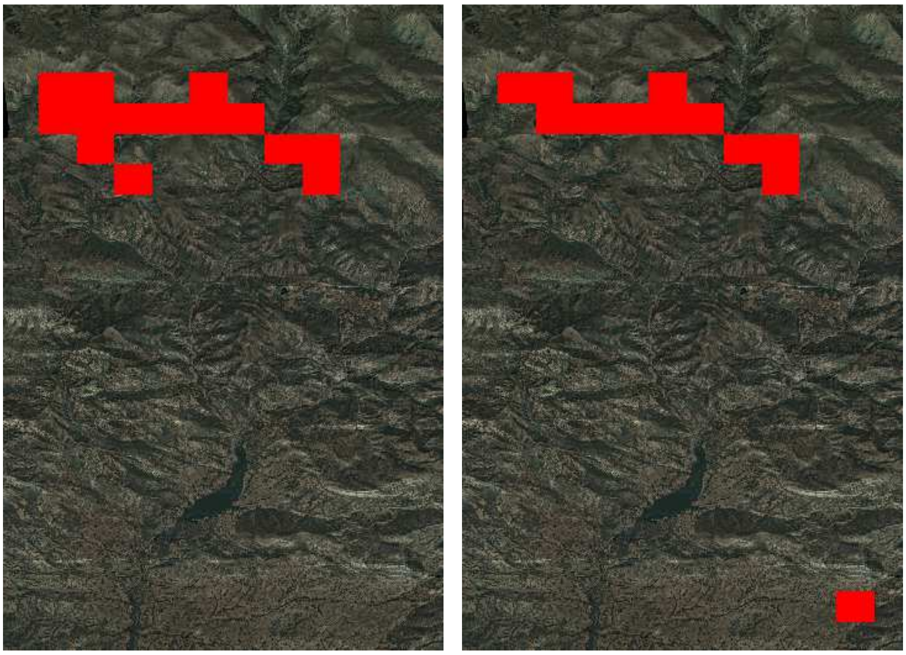

(c) Sampled every 0.03s at a grid size of 2000m.

(d) Sampled every 1.0s at a grid size of 2000m.

Figure 6.5: Examples of the view paths from a single person at different sampling frequencies and grid sizes. Each sub-figure has the same scale.

more frequently caused both the coverage and adjacency of traversed grid squares to increase, with coverage increasing steadily with sampling frequency whilst adjacency of traversed squares increases at a higher rate the more frequently the data is sampled. Increases in the grid size used to sample the data also caused steady increases in both coverage and square adjacency of the traversed path. Both the coverage and square adjacency of the position paths were slightly higher than those showing the view of the user.

There are a few interactions between the different design factors to note. Firstly, less frequent sampling did reduce the rate of coverage increase at lower grid sizes (up to ≈1000m grid size) and result in a slower and more consistent increase in adjacency as grid size increased, whilst higher sampling frequencies reached a high adjacency of approximately 3.5 at lower grid sizes which increased little beyond a grid size of 500m. The type of data sampled (position or view) had a larger effect on both coverage and adjacency at less frequent sampling rates.

6.2.4.2 Multiple People

The coverage and the adjacency of the grid squares for multiple users' paths under the different design factors is shown in figure 6.6, with examples of the effect these factors have on an example path shown in figures 6.7 and 6.8. In general, sampling frequency and the type of data used (position or view) had similar effects on the resultant coverage and traversed grid adjacency as with single user paths. However, both the coverage and grid adjacency of the paths increased with grid size in a different manner to single user paths, with each increasing at a higher rate at lower grid sizes before slowing down between approximately 500m and 1000m. The grid adjacency of position paths remained higher than those of paths depicting the view of users, but this was reversed for path coverage.

As with the single user paths there are some interactions to note. The first is that lower sampling frequencies again result in a more steady increase in adjacency as the grid size increases, but the difference between the most and least frequent sampling was notably lower than with a single user's paths. Secondly, there was a increased difference between the coverage of position and view data as the grid

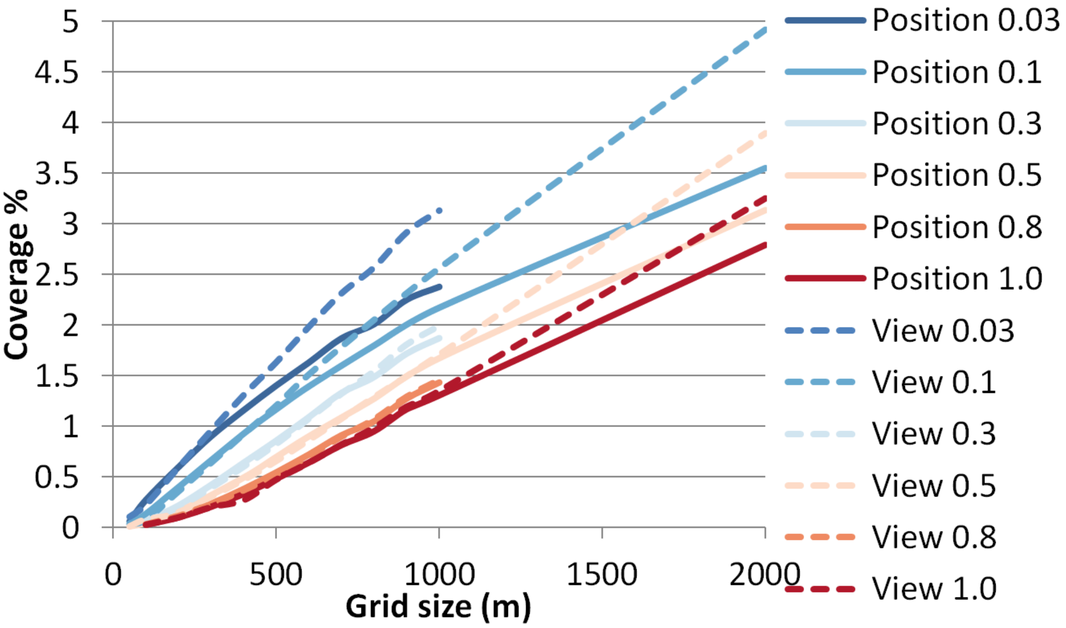

(a)

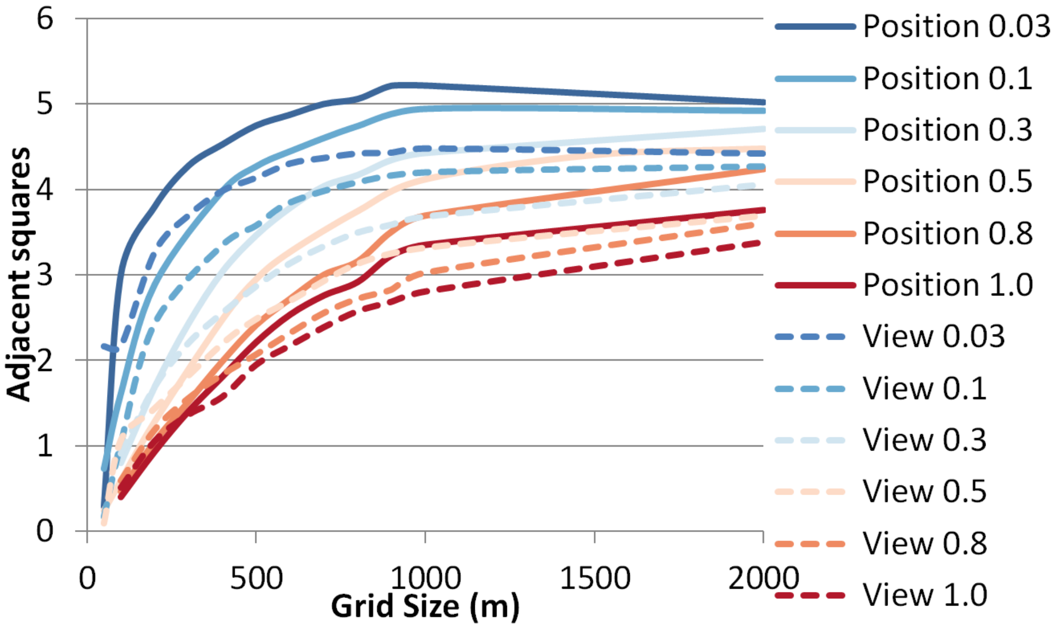

(b)

Figure 6.6: Results of analysing multiple people paths at different sampling intervals of time (between every frame (0.03s) and 1.0s) and grid size. The percentage of the environment covered by the path over a range of grid sizes is shown in (a). The average number of adjacent squares per populated square is shown in (b).

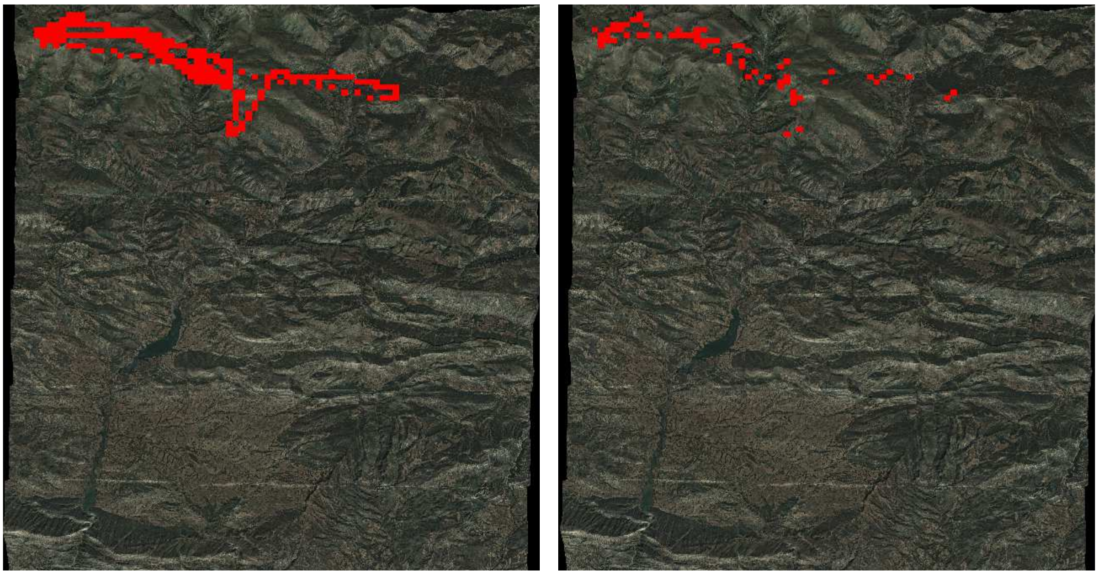

(a) Sampled every 0.03s at a grid size of 500m.

(b) Sampled every 1.0s at a grid size of 500m.

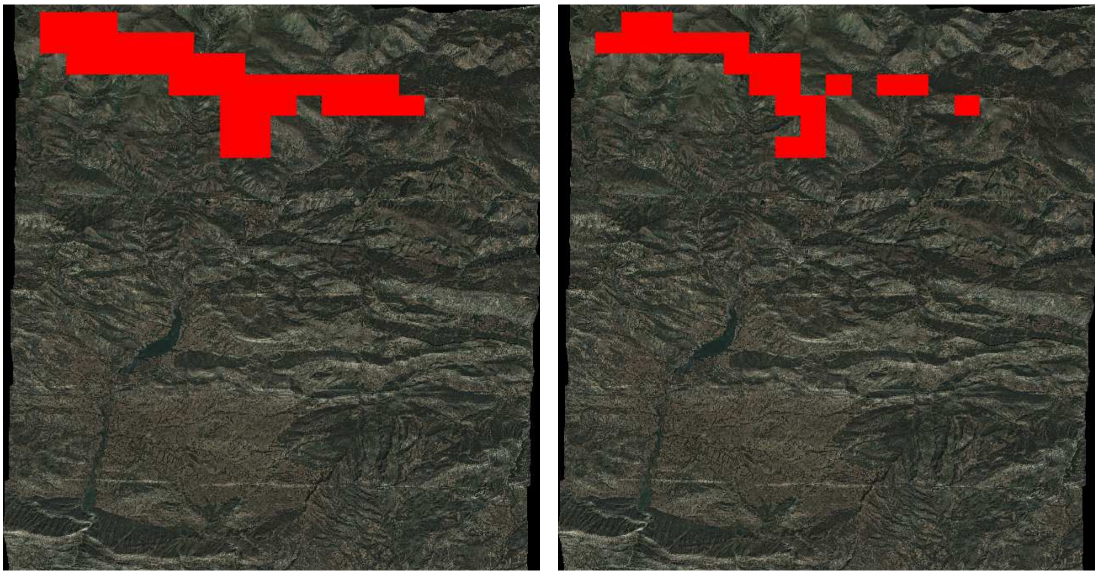

(c) Sampled every 0.03s at a grid size of 2000m.

(d) Sampled every 1.0s at a grid size of 2000m.

Figure 6.7: Examples of the position paths from multiple people at different sampling frequencies and grid sizes. Each sub-figure has the same scale.

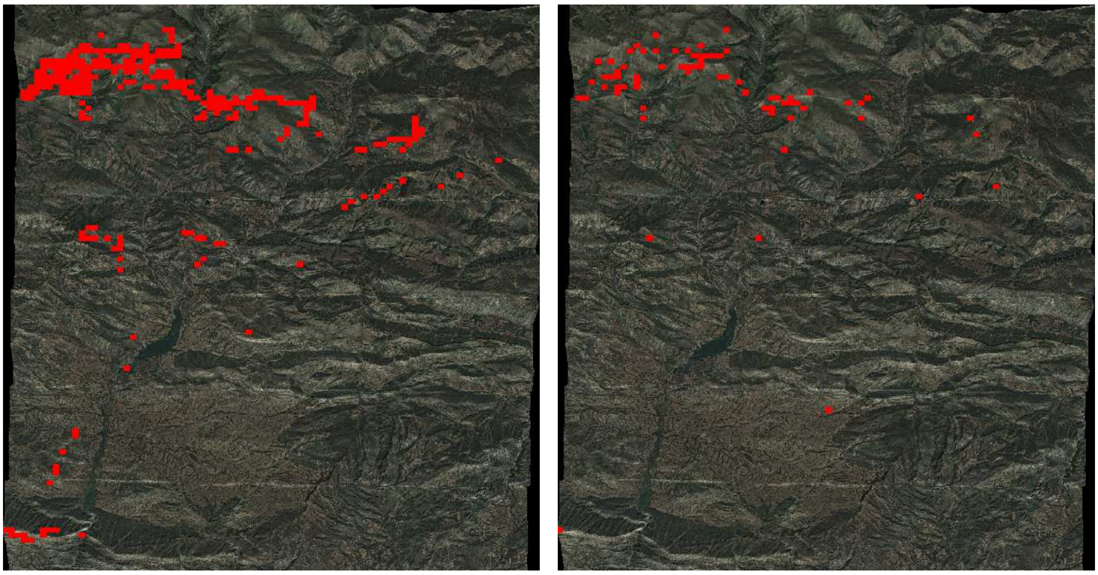

(a) Sampled every 0.03s at a grid size of 200m.

(b) Sampled every 1.0s at a grid size of 500m.

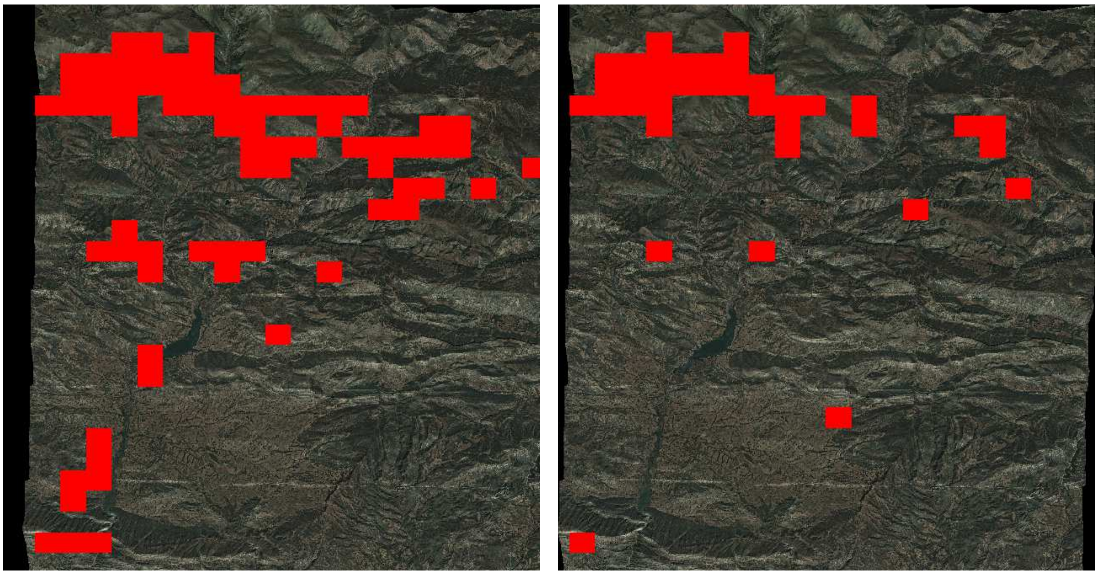

(c) Sampled every 0.03s at a grid size of 2000m.

(d) Sampled every 1.0s at a grid size of 2000m.

Figure 6.8: Examples of the view paths from multiple people at different sampling frequencies and grid sizes. Each sub-figure has the same scale.

size increases, with the coverage of view paths continuing to increase past ≈1000m whilst paths of position data increase at a lower rate past this grid size.

6.2.5 Discussion

Lower sampling rates caused the adjacency of the data to drop closer to the ideal value of 2, but do not provide a substantial reduction in the coverage of the data. Based on these results, the final path visualisations will use per frame data as there is little benefit in reducing the coverage of the path at the possible expense of obscuring the path details from the user. As the grid size of the data dropped below 500m, the adjacency of the cells dropped sharply. Considering that per frame data is being used, a grid size of 500m should be chosen to keep the data adjacency close to 2 whilst ensuring the data maintains a low level of coverage.

There was a notable increase in the adjacency of view data over position data in exchange for a small increase in coverage. Despite this, there isn't enough difference in the characteristics of the two data types to warrant investigating both, so only view data will be used for our path visualisation aid. The characteristics in the data were similar between single person and multiple people data and so the data for all of the participants captured will be used to give the best representation of usage with multiple users.

6.3 Second Iteration Designs

The type of visualisation used to represent the data is a factor that will be analysed and evaluated through a user study. The study will investigate what effect each representation has on the participants' performance and behaviour during navigation. Analysis of previous path visualisations (see tables 2.2, 2.3 and 2.4) determined that the three main forms of representation that could be used are ball of string, heatmap and markers (see section 2.4.2.2). The marker representation will be discounted from this experiment as the use of per frame data without the ability for the master to conduct on demand placement of markers will produce paths that are close in appearance to those produced by the ball of string representation.

A string representation allows users to follow a specific path, but the high number of populated adjacent squares from data of multiple users can cause crossings of these paths as discussed in section 2.4.2. One way of reducing this clutter is to use the data of previous users to calculate the most popular path between the targets and only show this path to participants. To calculate this each grid square will be used as a node of a graph, with a unidirectional link between each node and the surrounding ones representing traversal between each grid square. The cost of traversing each link is defined as the inverse of the frequency of traversal between the regions multiplied by an arbitrarily large number to get a integer cost. A shortest path algorithm [40] can then be used to calculate the most frequently traversed path between two points (e.g., the starting point and a target).

This path calculation is used to cost the paths between the starting position and each of the three targets, and then the paths between each of the three targets. The total cost of visiting each of the targets in different orders is determined by adding the costs of the components so the lowest costing (and therefore most frequently traversed) overall path can be selected.

The coverage and adjacency of the popular paths are presented in figure 6.9. The paths resulting from this method have coverage and square adjacency which is comparable to paths which sampled the data at frequencies of 0.5 and 1.0 seconds. The lower coverage allows a larger area to be dismissed from the user's search, whilst the average adjacency not rising above 3 demonstrates that there are few branches, crossing or other complications to these paths. Examples of string paths using the original and most popular paths are shown in figures 6.10 and 6.11.

For the heatmap representation the manner in which the colour range is mapped can be scaled according to the data it is representing. The maximum value of the data used in this scenario quickly increases as more users' movements are added, but these maximum values are only present in a small number of sample squares, with many squares only being briefly and infrequently visited. If this was linearly mapped onto the colour range then it would result in a visualisation with only a few visible areas (see figure 6.12).

Instead, with the colour value being used on a 0 to 1 scale, any sample regions which have been visited will be increased by 0.3. This will result in areas that have only been visited infrequently being visible, whilst the small number of high

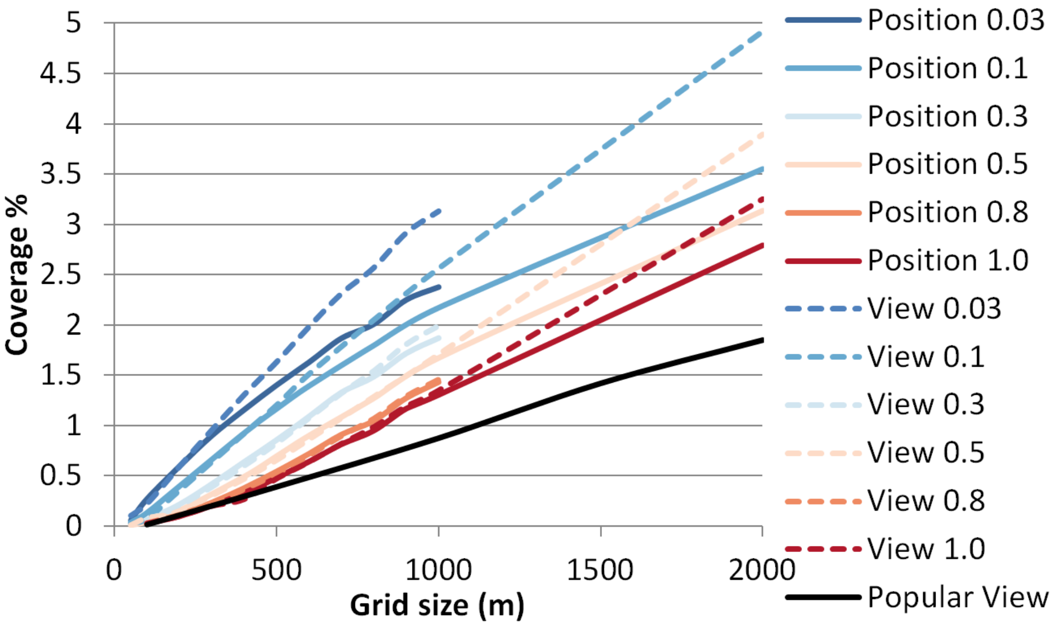

(a)

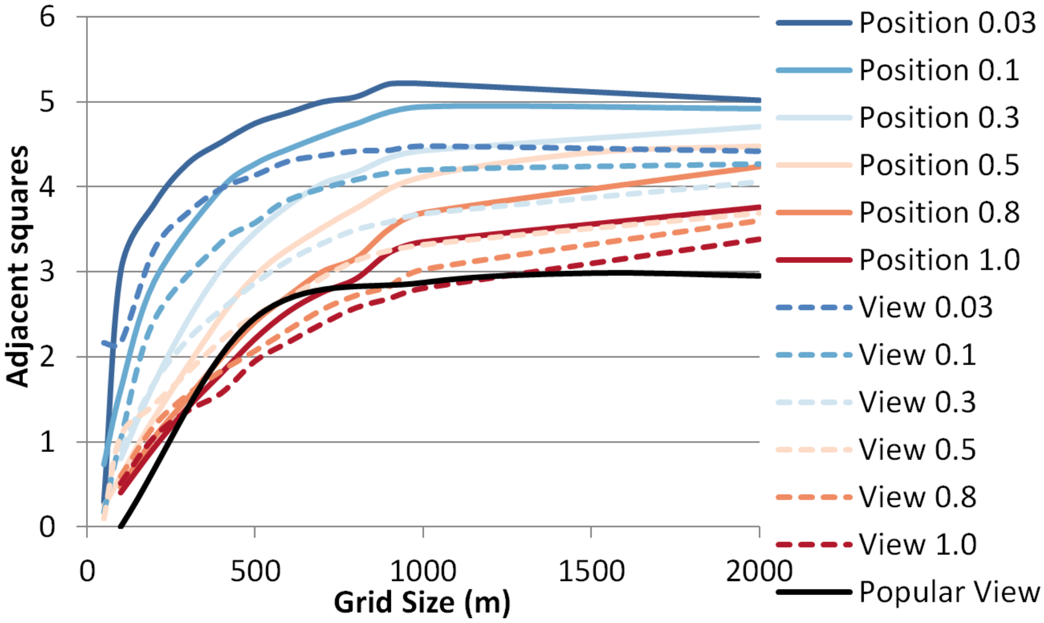

(b)

Figure 6.9: Results of analysing the paths for multiple people with a subset of the original data from figure 6.6 and the most popular traversed paths. The percentage of the environment covered by the path over a range of grid sizes is shown in (a). The average number of adjacent squares per populated square is shown in (b).

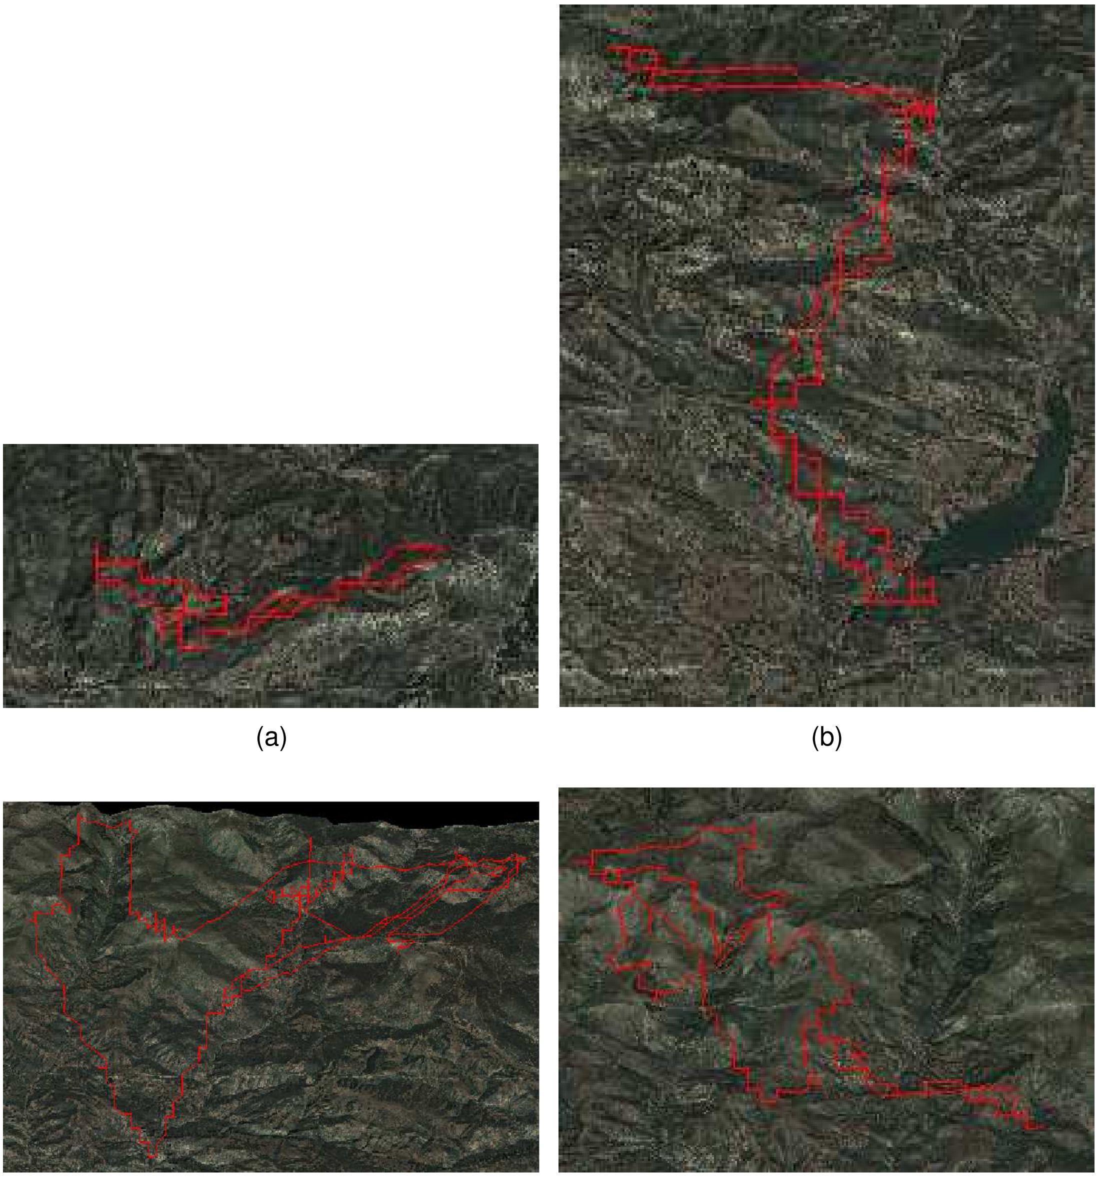

(a) (b)

(c) (d)

Figure 6.10: Examples of paths the four wall-sized master users took between targets during the trials.

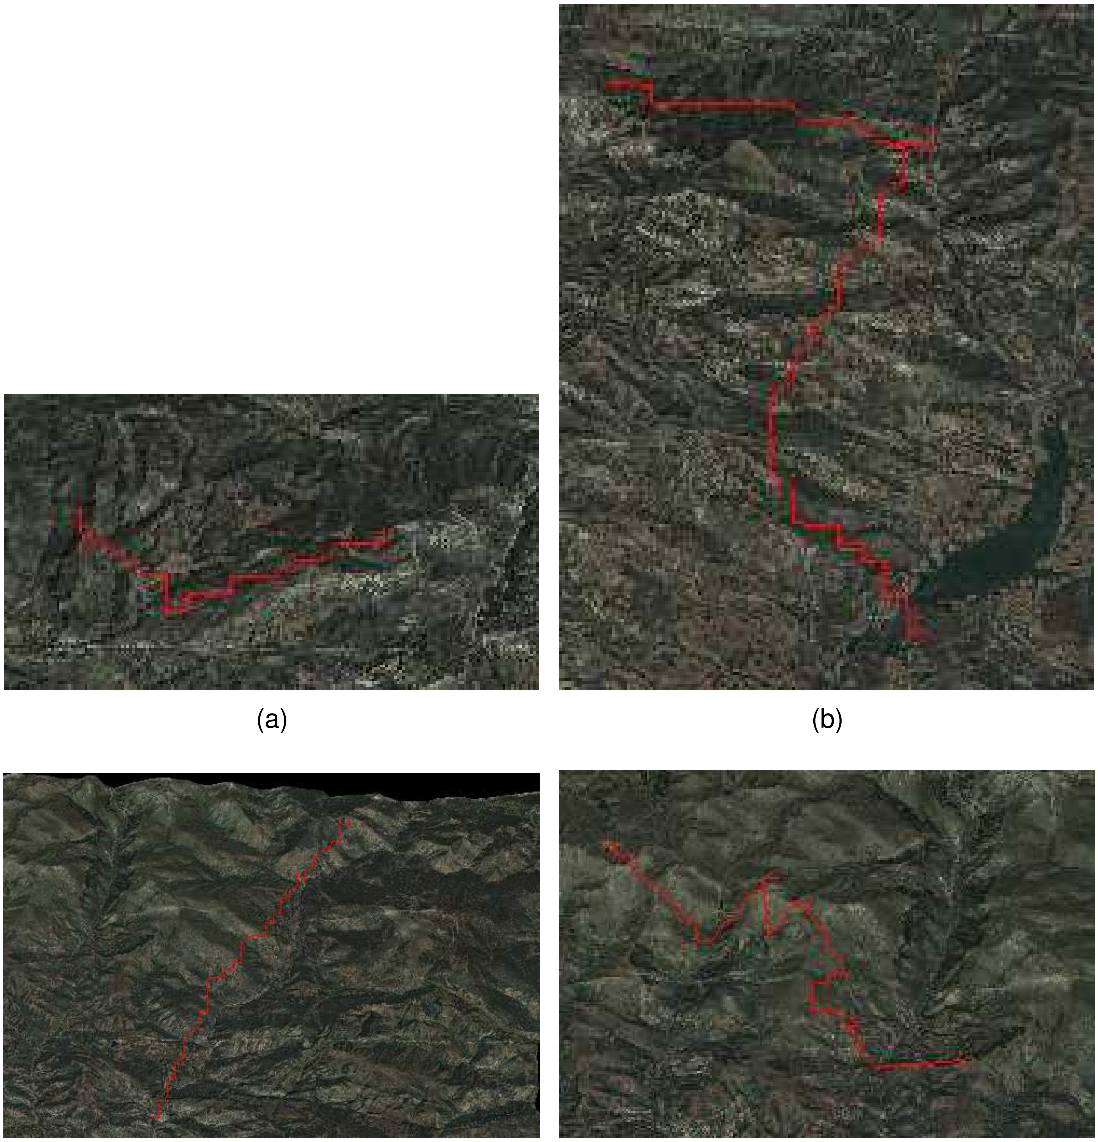

(a) (b)

(c) (d)

Figure 6.11: Examples of the string path visualisation provided to participants during the trials.

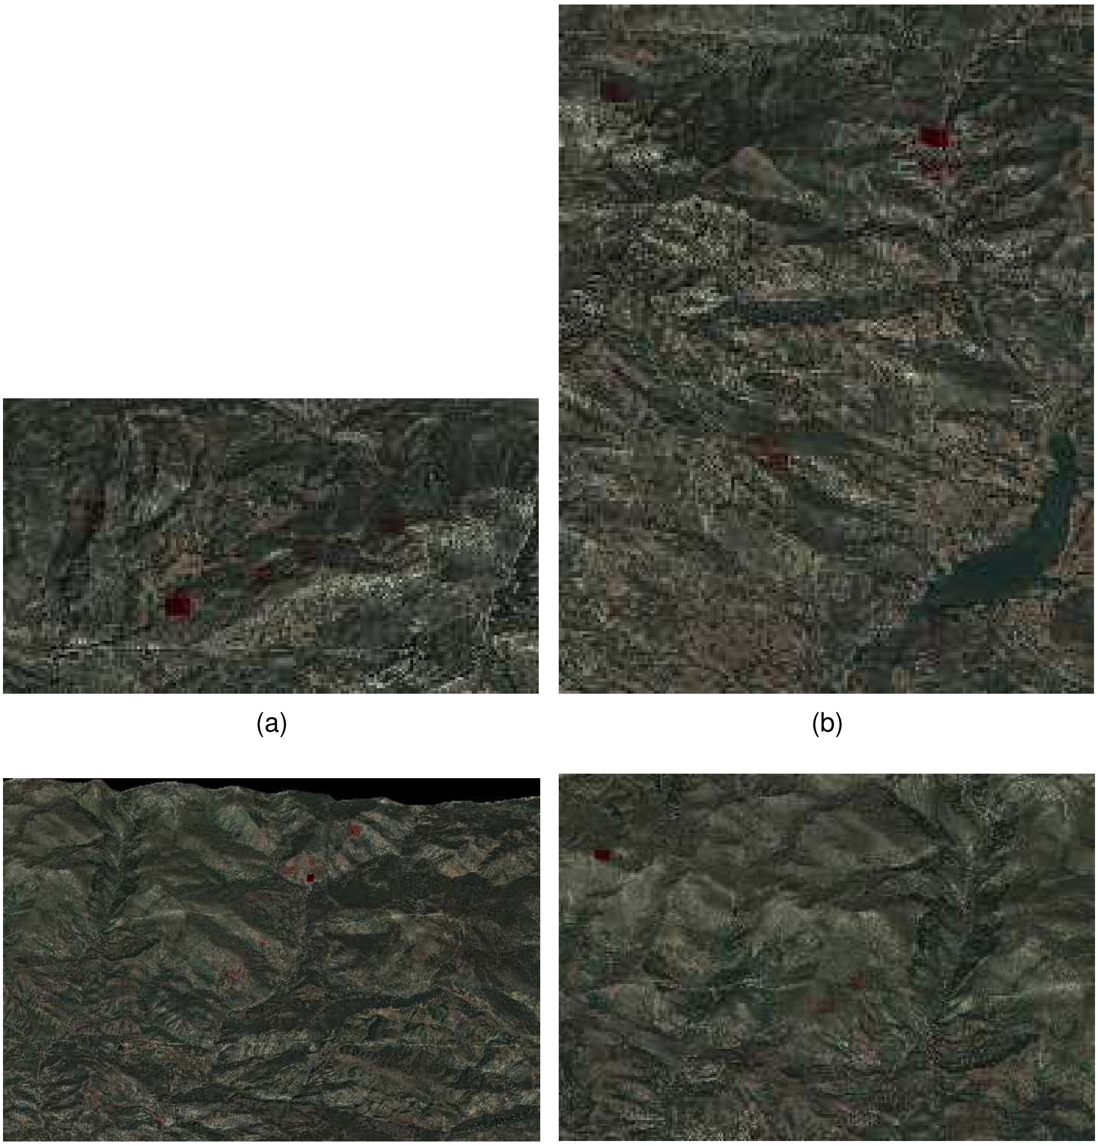

(a) (b)

(c) (d)

Figure 6.12: Examples of the heatmap path visualisation with frequency being linearly mapped to colour intensity.

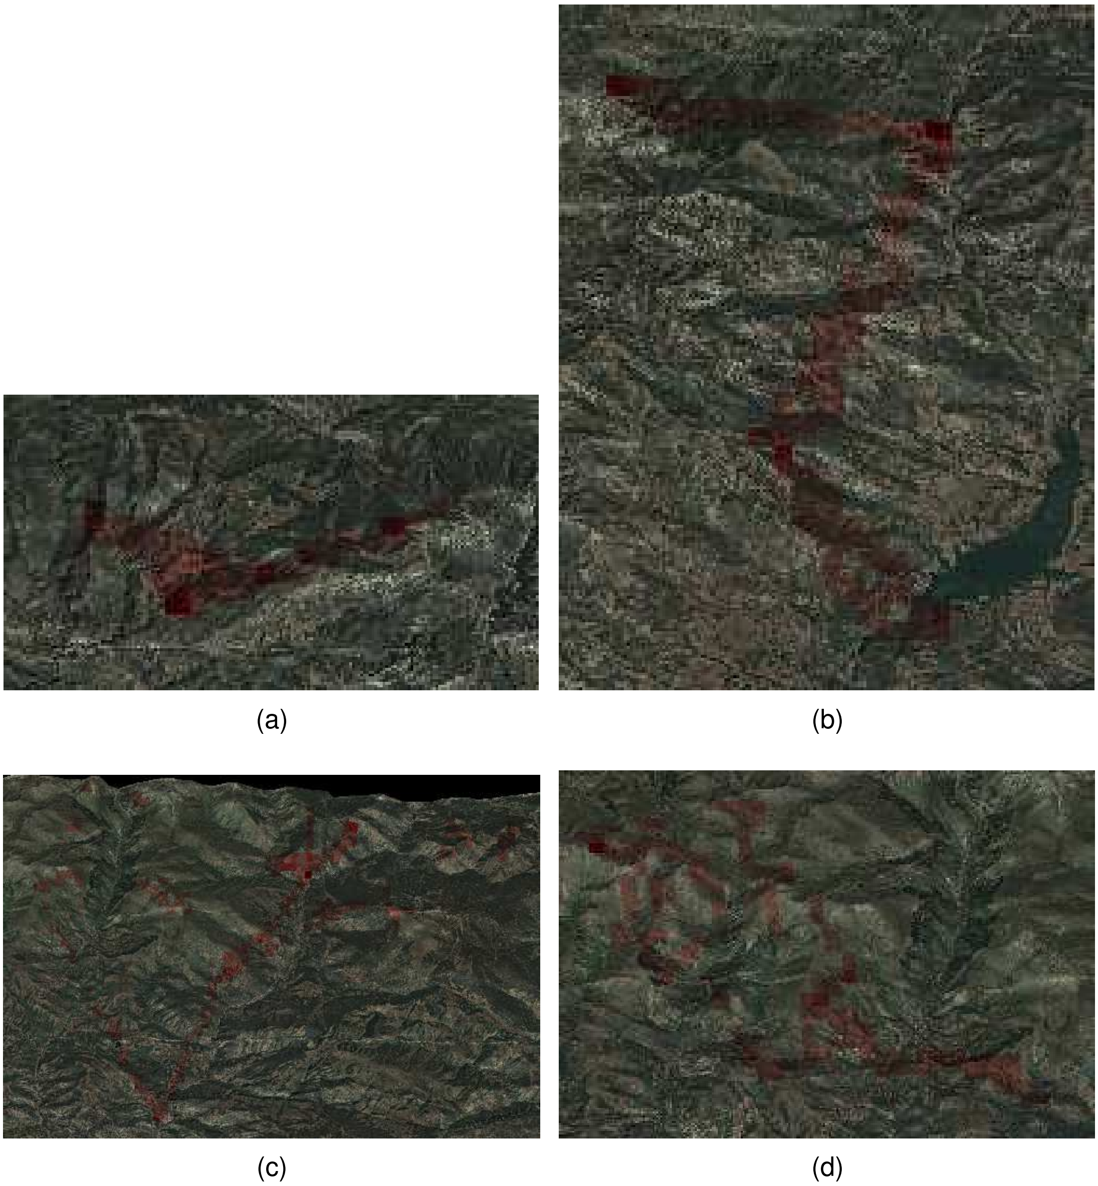

(a) (b) (c) (d)

Figure 6.13: Examples of the heatmap path visualisation provided to participants during the trials.

frequency regions will ensure that the compression of the highest 30% of potential frequency onto the same colour does not misrepresent the data to a large degree.

One issue with this is that the maximum colour value will be the same over a range of high frequencies but there are usually only a few regions with a high frequency of visits which do not need to be accurately distinguished. This results in a heatmap that is fully populated and should allow users to understand the range of movements used by previous users rather than just a small proportion of the most visited areas. Examples of the resulting heatmaps are shown in figure 6.13.

6.4 Experiment 4: Navigation when using path visualisations

This experiment investigated the effect the path visualisations developed in section 6.3 had on the navigation errors of users in the master slave scenario used in previous experiments. Our hypotheses were as follows:

H4.1 Both path visualisation aids would result in fewer task failures than the base condition.

H4.2 The string visualisation would allow participants to find and traverse a navigable path between targets more quickly that the heatmap condition.

6.4.1 Method

6.4.1.1 Factors

The experiment had one within-participant factor:

- Type of path visualisation used (Base (none) vs. String vs. Heatmap)

This factor allows both H4.1 and H4.2 to be tested.

6.4.1.2 Measures

The following measures were taken during each trial:

- Success/ failure in finding each target
- Time elapsed during the trial
- Visibility of the path visualisation per frame (in String and Heatmap conditions)
- Position and orientation per frame

The first two measures allowed performance to be assessed, the next factor allowed analysis of how frequently the path visualisations were used while the last measure allowed participants' navigational behaviour to be investigated. Instances of common ground breakdown were observable through reduced navigational performance and/ or differences in navigational behaviour.

6.4.1.3 Participants

Twelve slave participants (four females and eight males, mean age = 26.5 years, *SD* = 4.9) and the same four master users as the previous experiment (one female and three males, mean age = 24.0 years, *SD* = 2.2) took part. The type of path visualisation was viewed by participants in a Latin Square arrangement.

6.4.1.4 Materials

The materials for this experiment consisted of the same geophysics dataset, wall-sized master system to record the paths and desktop slave system used by participants as Experiment 3. The slave participants could toggle whether the path visualisation (if applicable) was visible by pressing a button on the gamepad. The visualisation was not visible at the start of each trial.

6.4.1.5 Procedure - Master

The same master paths sets were used as in Experiment 3 with the path visualisations being created from the data of all four master users for each trial.

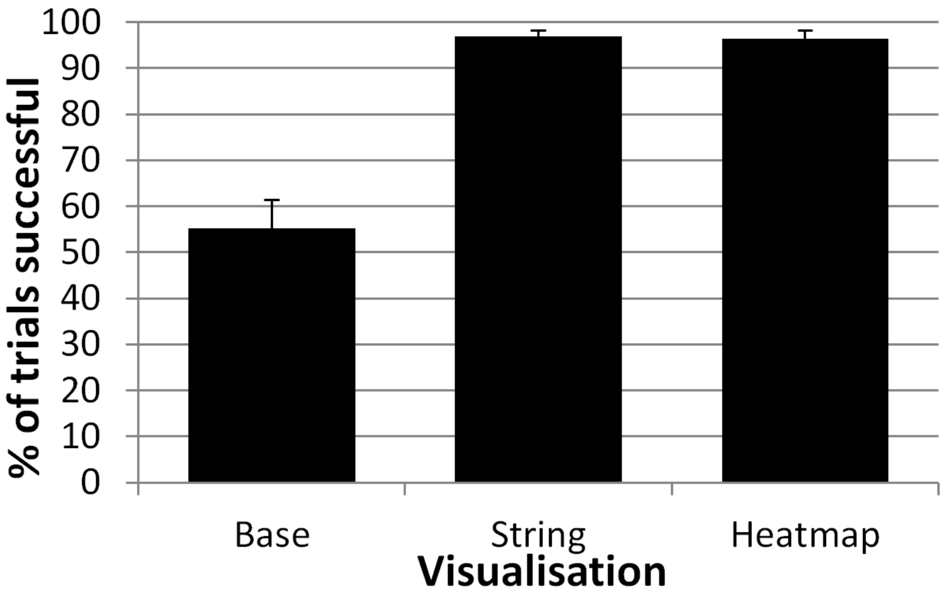

Figure 6.14: Percentage of trials that were successful for each visualisation condition. Error bars show standard error of the mean.

6.4.1.6 Procedure - Slave

The procedure for the slave participants was the same as in Experiment 3 with the path visualisation condition replacing the type of view used. In each trial, the time and movements that participants took, alongside whether the path visualisation was toggled on or off, were recorded for subsequent analysis.

6.4.2 Results

This section first analyses the general performance of participants at the task, before identifying the major behaviours that contribute to any differences in performance.

6.4.2.1 Performance

The visualisations provided were expected to result in changes to the quantity of failures that participants exhibited, so the percentage of successful trials was analysed using a repeated measures ANOVA that treated the visualisation used (base

View	Experiment 3			Path Visualisation	Experiment 4		
	Not Close	Close	Total		Not Close	Close	Total
Base	16.7	24.5	41.1	Base	26.6	18.2	44.8
Context	27.6	16.7	44.3	String	1.6	1.6	3.1
Local	20.8	21.4	42.2	Heatmap	1.0	2.6	3.6

Table 6.2: Percentage of trials of each type of failure under different conditions in Experiment 3 and 4.

vs. string vs. heatmap) as a within participants factor. This analysis showed that there was a significant effect of the visualisation (*F*(2, 10) = 33.61, *p* $<$.01) on success (see figure 6.14). Means varied from 55% (*SD* = 22.2) for the base visualisation to 97% and 96% (*SD* = 4.2 and 6.2) for the string and heatmap visualisations, respectively. Post hoc, pair-wise comparisons performed using the marginal means showed that participants were significantly more successful in the string and heatmap condition than in the base condition (p $<$.01 in both cases), but performance between the string and heatmap conditions was similar (p $>$.05).

Success in finding additional targets in the base condition decreased between the first (*M* = 95%, *SD* = 7.0), second (*M* = 81%, *SD* = 15.2) and third target (*M* = 55%, *SD* = 21.3) in a similar manner to previous experiments. The types of failures that occurred were classified in the same way as previous experiments (see section 4.3.2.3 and 5.3.2.1) and also resulted in similar classification patterns as the previous study (see table 6.2).

There was also a significant difference between the different visualisation conditions in the time taken to find the targets when participants were successful (*F*(2, 10) = 5.52, *p* $<$.05) (see figure 6.15). The means varied from 75.3 seconds (*SD* = 23.0) for the base condition to 52.1 seconds (*SD* = 21.7) for the string visualisation and 42.3 seconds (*SD* = 24.4) for the heatmap visualisation. Post hoc, pair-wise comparisons performed using the marginal means showed that participants completed their task significantly faster in the string and heatmap conditions than in the base condition (p $<$.05 for string and p $<$.01 for heatmap), and significantly quicker in the heatmap condition than in the string condition (p $<$.05).

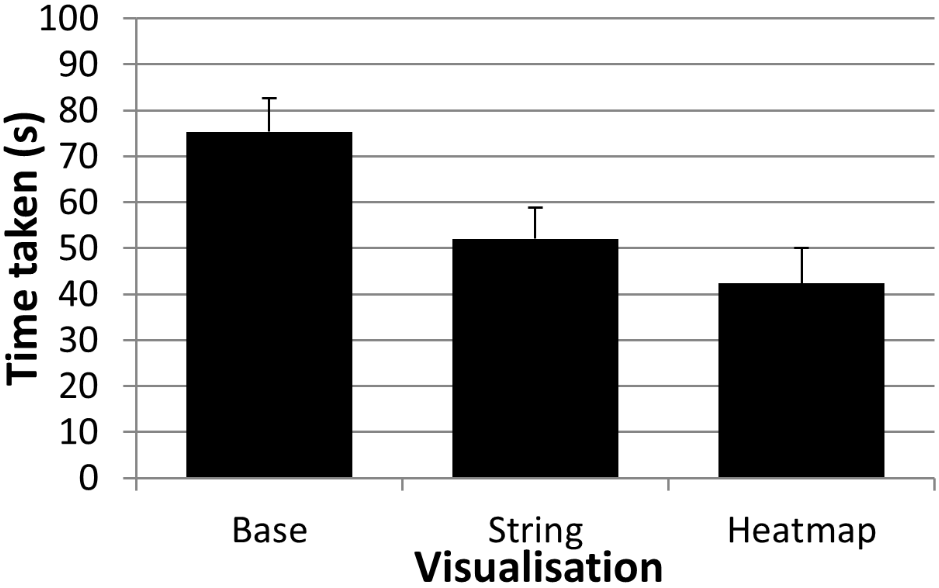

Figure 6.15: Average time taken to successfully complete trials for each visualisation condition. Error bars show standard error of the mean.

There was no significant difference in the percentage of time the visualisations were visible while conducting the trials (76.6% for string vs 74.7% for heatmap) but there was a large variation between participants in the usage (a SD of 30.5 for string and 31.42 for heatmap). There is a fairly clear division in participants either displaying the visualisations almost constantly ($>$80 % of the time), or only at specific points in time ($<$40% of the time), with most participants being in the first group in both conditions (8 out of 12 participants).

6.4.2.2 Behaviour

The string and heatmap conditions provided navigational information that resulted in both superior performance and some key behavioural changes. In the base condition (as in previous previous experiments) some participants started trials by spending long periods of time in the area which did not contain the targets (see figure 6.16). This behaviour was absent within both string and heatmap visualisations. When participants did manage to establish areas in which they thought targets were located in the base condition, this could still be inaccurate and require them to tra-

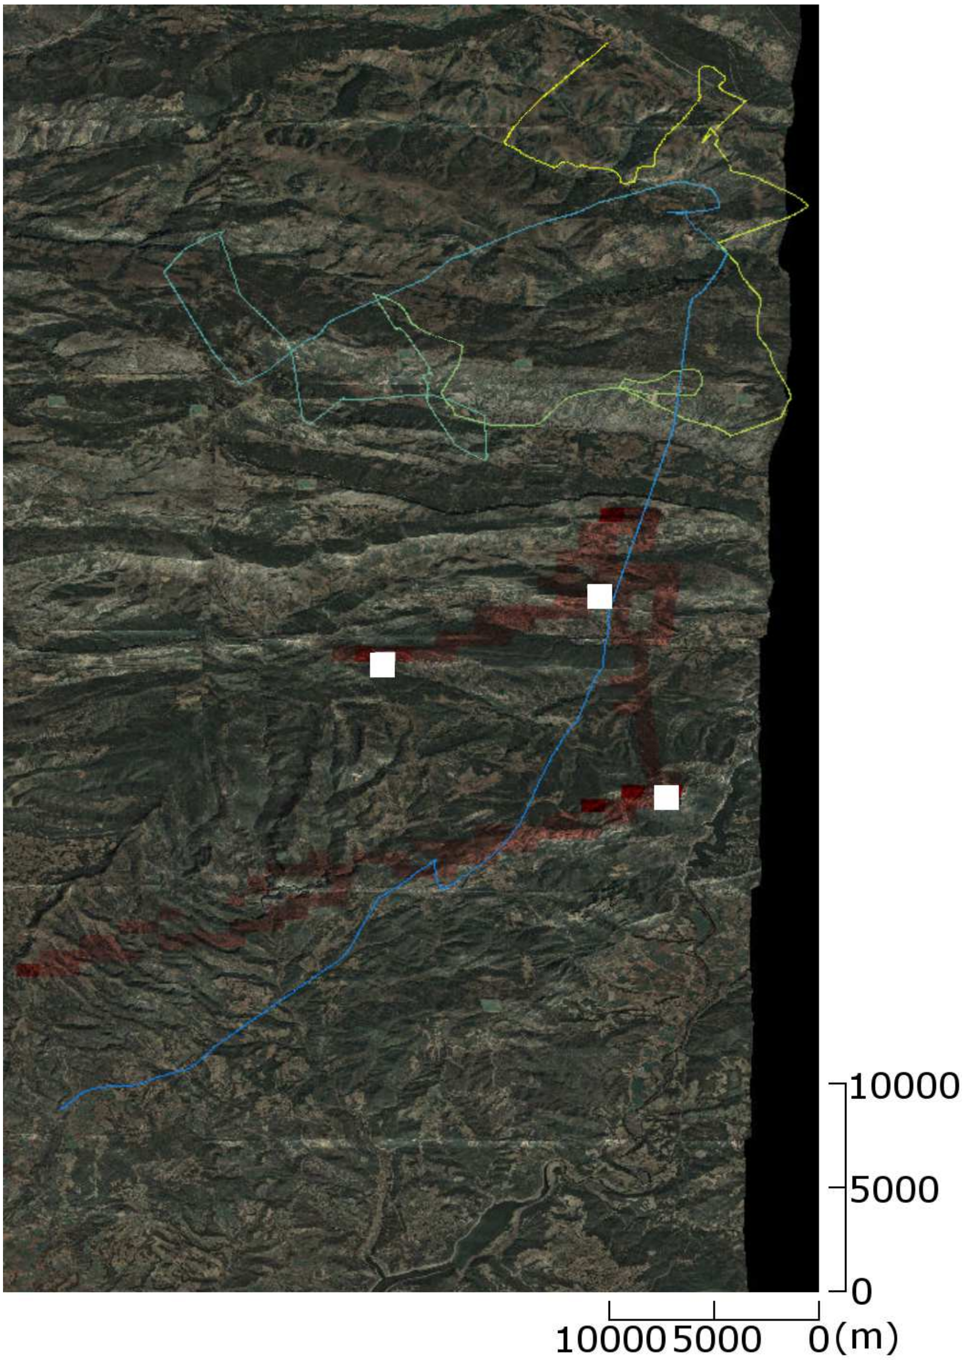

Figure 6.16: Example of a participant in the base condition spending most of their time in an area away from where the targets and previous paths are located. The participant's path is shown as a line going from blue to yellow as time progresses, the target points highlighted in white and the heatmap visualisation that would be shown for this trial shown in red.

verse this area carefully through extensive navigation or repeated re-visitation of the area (see figure 6.17a). The use of the string and heatmap visualisations allowed participants to focus on a smaller area or reduce the number of repeated visitations (see figure 6.17b).

With the string and heatmap conditions participants adopted other particular behaviours. Participants frequently followed the same, or similar route to the one that the string path showed them between the different targets, a situation in which the path simplified their navigation task (see figure 6.18). However, there were other occasions when participants wanted to follow a different path to the one provided by the string visualisation, e.g., potentially wanting to follow the original path they viewed or taking a shortcut to a target that was close. Either of these situations required them to traverse a new path rather than an existing path, potentially reducing the speed with which the task will be performed. It is in these scenarios where the heatmap condition may be more beneficial than the string condition to the participant, as it is more likely to show a potential path that is a close match to their proposed movement (see figure 6.19). Although the larger number of potential paths requires a decision of which path to use, following this path reduces the difficulty of navigation when they have started movement.

This behaviour was investigated by calculating the percentage of time during each trial the participant was in the proximity of the two representations. This was calculated through testing whether each point along the participant's path was within the same grid square as a point of the string representation or square of the heatmap representation. This was analysed using a repeated measures ANOVA that treated the visualisation used (base vs. string vs. heatmap) and representation covered (string vs. heatmap) as within participants factors. This analysis showed that participants spent significantly more of trials on the heatmap representation (M = 36% of the trial) than the string representation (M = 21% of the trial, $F(1, 11) = 530.73$, $p < .01$) (see figure 6.20).

There was also a significant effect of the aid used ($F(2, 10) = 16.48$, $p < .01$) with participants spending significantly less time near the representations in the base condition than either of the string or heatmap conditions (M = 18% vs. 35% vs. 32%). Post hoc, pair-wise comparisons performed using the marginal means

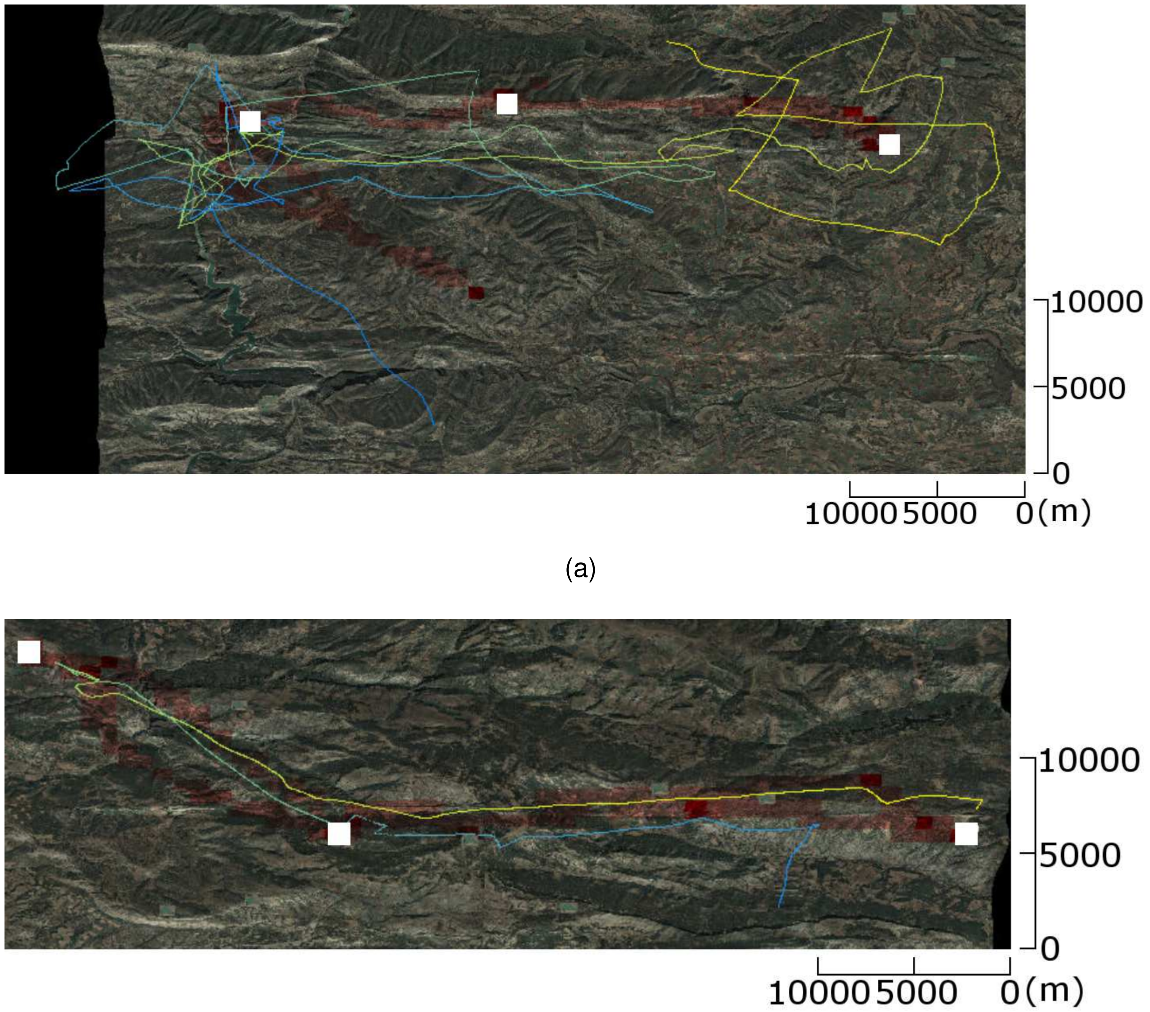

(a)

(b)

Figure 6.17: Example of participant behaviour when close to the target in base (a) and heatmap conditions (b). The participant's path is shown as a line going from blue to yellow as time progresses, the target points highlighted in white and the heatmap visualisation that would be shown for this trial shown in red.

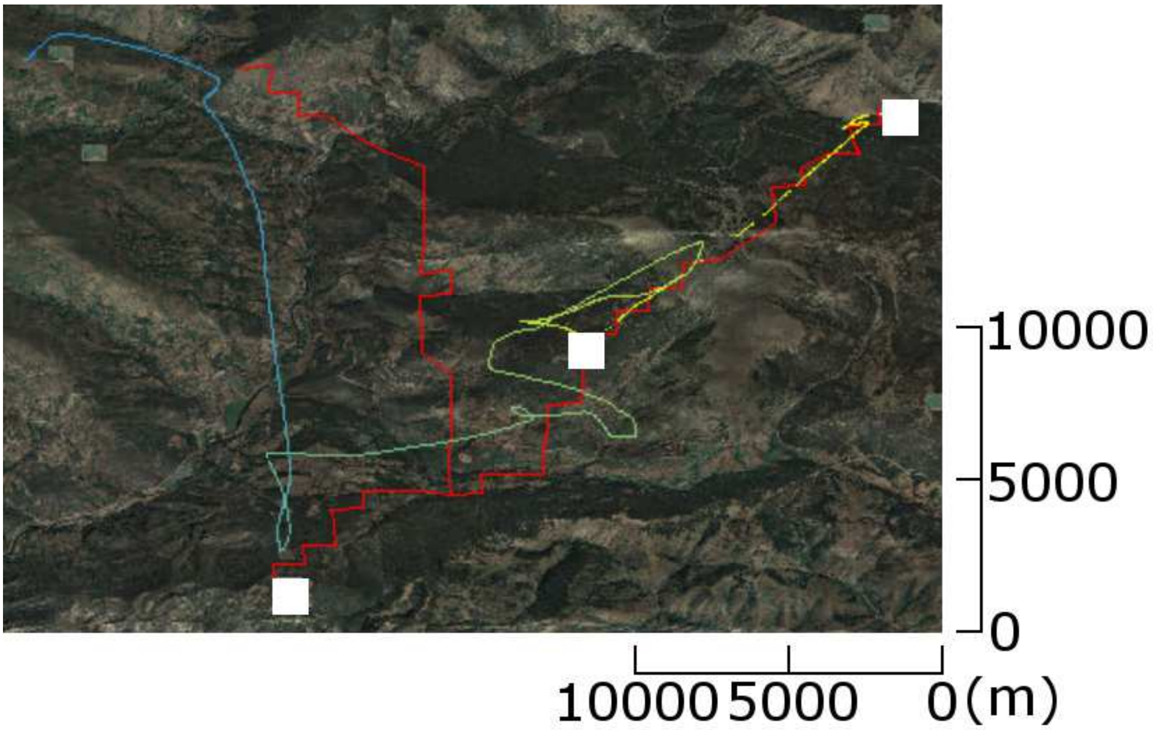

Figure 6.18: Example of participant following the string visualisation. The participant's path is shown as a line going from blue to yellow as time progresses with the string visualisation that was shown for this trial shown in red.

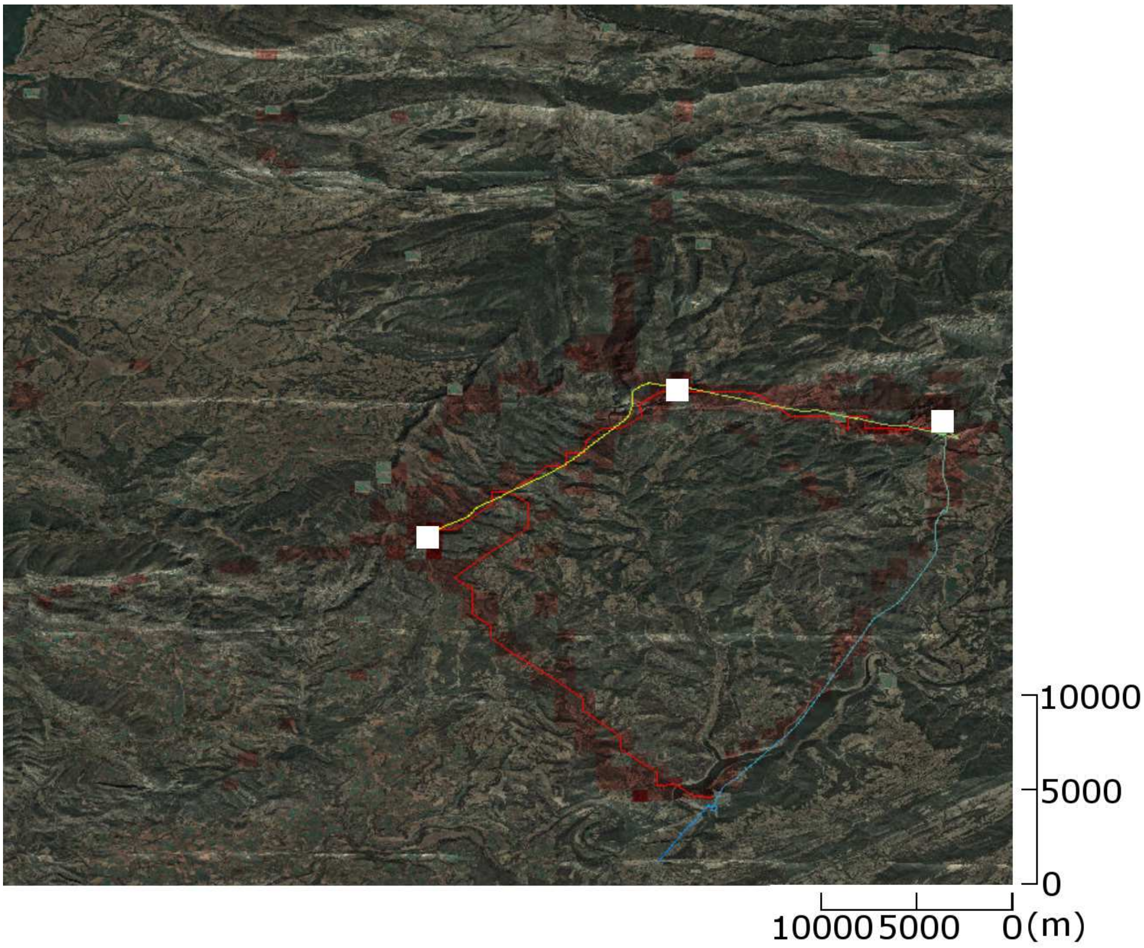

Figure 6.19: Example of participant behaviour when following the heatmap visualisation with the heatmap visualisation visible. The participant's path is shown as a line going from blue to yellow as time progresses. The trial's string visualisation is also shown in red.

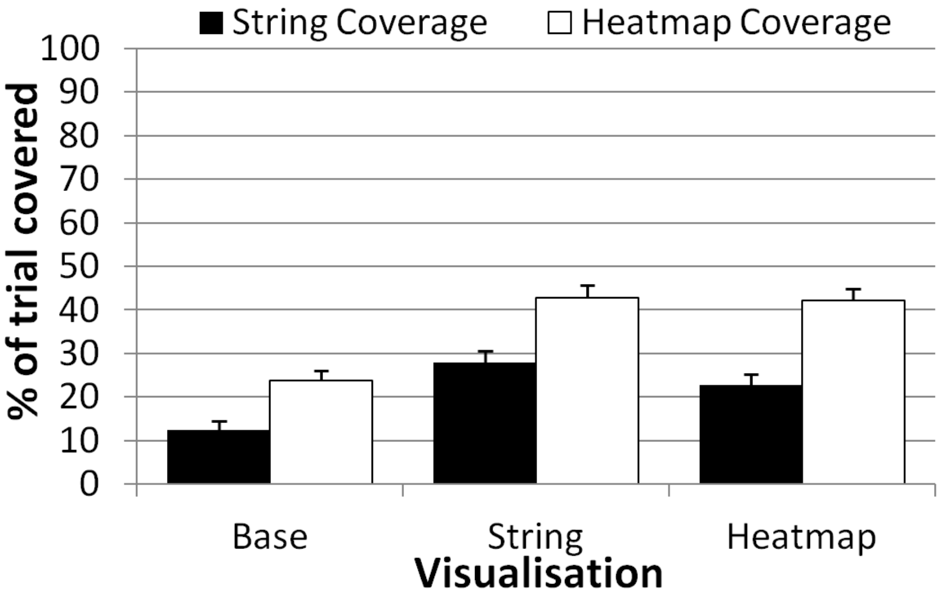

Figure 6.20: Average percentage of the trials that covered the representation for each visualisation used. Error bars show standard error of the mean.

showed that participants were closer to the path for a significantly greater percentage of the trial in the string and heatmap conditions than in the base condition (*p* $<$.01 for both string and heatmap), but there was no significant difference between the heatmap condition and string condition (*p* $>$.05).

In addition there was a significant interaction between the type of representation used and the coverage of each representation during the trial (*F*(2, 10) = 14.294, *p* $<$.01). Participants were close to the heatmap representation an approximately equal percentage of the trial with heatmap and string representations (*M* = 43% vs. 42%) but participants were close to the string path for a greater percentage of the trials when the string representation was used than when using the heatmap representation (*M* = 28% vs. 23%).

6.4.3 Discussion

As hypothesised (see H4.1) the use of path visualisations results in a clear reduction to both the failure rate and time taken to find the targets due to participants' behaviour exhibiting better search patterns than in the base condition. However, contrary to H4.2, participants were able to use the heatmap visualisations to complete the task in less time than the string condition. The differences in information provided by each of the visualisations affected how participants could use them to follow a previous path which may have contributed to this difference.

There are other potential reasons behind the performance difference of the string and heatmap visualisations which were not analysed. One reason could be that although the string data emphasises areas of the environment that were frequently traversed and may be navigationally important at a contextual scale, these may not be the areas which users view in greater detail. In more complex data analysis tasks this difference may have further effects, due to it not indicating which areas have been extensively analysed and which have not.

Secondly, although the string path shows a specific path that could be followed by participants, it does not allow users to immediately exclude areas outside of this path that have not been visited from their search. However, the heatmap visualisation provides information on areas that have been visited at least once and so presents an immediate overview of those areas have been visited and those that have not.

Lastly, an unexpected consequence of the heatmap design that may contribute to some of the performance difference is that the green targets could be more easily seen when the red heatmap was placed underneath them than the normal landscape colouring which existed with the string visualisation.

6.5 Summary

This chapter investigated the effect that path visualisations had on the common ground breakdowns within wall-sized master to desktop slave collaborative navigation. The variables affecting the design of such an aid were investigated as part of a technical evaluation before a experiment evaluated two different representations of previous paths.

Experiment 4 evaluated the effect of two path visualisations, string and heatmap, on participant performance compared to the base condition (H4.1) and between the two path visualisation conditions (H4.2). Results show that both representations significantly increased participants' success rate and reduced the time taken to complete the task, with the time taken also being significantly lower in the heatmap condition than the string condition. Behavioural data showed that participants visited areas unrelated to their search less frequently and remained closer to the desired path in both heatmap and string conditions.

Chapter 7

Conclusion and Future Work

This chapter concludes the thesis by reviewing and assessing the major contributions it provides before discussing some other notable findings. The chapter ends by discussing potential future work that could expand upon the current research.

7.1 Conclusion

The research within this thesis sought to investigate and reduce the frequency and severity of common ground breakdowns between desktop and wall-sized display users in a distributed scenario. This has led to four major contributions which are summarised and assessed in sections 7.1.1 to 7.1.4. Some other notable findings that can be taken from this research are described in section 7.1.5.

Previous work related to the topics within this thesis are presented in chapter 2, including a introduction of common ground (and the associated processes of grounding and breakdown), CSCW classifications, navigation strategies for real and virtual worlds, CVE navigation aids and multiple display arrangements.

A suitable scenario in which collaborative, master/slave navigation across wall-sized and desktop displays could occur is derived from an existing geological fieldtrip in section 2.6. Section 2.7 explores the forms of common ground breakdown that occur within this scenario and presents examples of functionality that have been previously successful in reducing similar forms of breakdown. The platform that was developed to conduct the subsequent experiments is discussed in chapter 3.

7.1.1 Contribution 1 - Classification of breakdowns between wall-sized and desktop display users

Chapter 4 presents two experiments that classify the frequency, severity and conditions of common ground breakdown, as well as assessing examples from a range of standard navigational factors to guide later research. Initial hypotheses for Experiment 1 stated that increased view similarity through a wide FOV context view (H1.1) and less complex navigational tasks (H1.2) should lead to reduced common ground breakdown. Both of these hypotheses were found to be false, likely due to the limited nature of the task (re-visiting a single target after viewing the navigational actions of a wall-sized master user).

Experiment 2 therefore hypothesised that common ground breakdown would be more prevalent when the master control interface was different to the participants' (H2.1), task complexity was increased (H2.2) or the participant was interrupted (H2.3). The results showed H2.1 to be false, but supported both H2.2 and H2.3. Throughout the two experiments participants exhibited a high level of breakdown across different view, movement and master input device conditions, resulting in participants failing to complete 45% of trials with multiple targets (see section 4.3.2). Despite the level of failures, participants frequently managed to get close to a target they were looking for and attempted to follow the original path in many instances (see figure 4.7).

Although the specifics of this contribution are subject to various design choices that were made for this specific scenario and experimental platform (see section 2.6.2), many of the findings can be applied to other scenarios involving analysis of large, multi-scale datasets through wall-sized and desktop display collaboration. Relevant elements of this scenario that may have impacted upon the types and quantity of breakdowns (e.g., task types, control interface and environment type) have been clearly stated so that predictions can be made when applying findings from this work to other scenarios. However, it is important to note that results from these experiments indicated that some standard navigational variables (e.g., master interface) had little effect on overall navigation performance and breakdowns.

One such element that can be considered in this manner is the finding that in a notable amount of failures, participants were able to get close to the target they were searching for but not recognise the target whilst there. In some cases the type of environment may reduce the absolute quantity of these spatial breakdowns, for example, navigation of environments with a clear multi-scale structure [79] or high number of local landmarks [107] should result in fewer breakdowns.

However, the combination of landmarks in this scenario is similar to those that would be experienced in other low fidelity visual environments such as abstract data spaces, allowing such findings to be used in those environments and require the use of additional navigational aids to pre-empt these breakdowns (e.g., adding artificial landmarks [92]). These findings would also be effective for the substantial minority of users who also experience difficulty navigating more detailed environments despite the higher number of salient landmarks present [19, 76, 85].

This classification of breakdowns can also be used to inform extensions of the conditions found within this scenario. Master/slave navigation was used in this scenario, but it is likely that users will want to control their own navigation around the environment at certain times in longer and more complex collaborations. Under these conditions the common ground between users will be further lowered and spatial orientation and data visibility breakdown likely to be more frequent even in homogeneous desktop collaboration [42].

7.1.2 Contribution 2 - No effect of additional views on breakdowns

Chapter 5 investigated whether two examples of additional views, a navigational aid that has been previously successful in homogeneous desktop navigation, would aid participants within this scenario. Experiment 3 hypothesised that a wide FOV context view and a local map view would improve participants route (H3.1) and survey knowledge respectively (H3.2) and that participants would be more successful when using the same interface as the master (H3.3). However, the results indicated that neither view condition had an effect on the frequency with which participants successfully completed the task (see figure 5.3), showing that all three hypotheses to be false. Use of the context view increased the severity of the failures, which further analysis attributed to behavioural differences such as attempting to complete a route to an incorrect target (see figure 5.9) and using the context view to view the environment in detail and context simultaneously (see figure 5.10).

The difficulty participants had in integrating additional views of the environment together results in these forms of aids not being recommended for other work under similar conditions. However, it is unclear to what degree these difficulties can be attributed to master/ slave navigation being used, the differences from heterogeneous display use or the combination of these two factors.

This form of aid may still be beneficial in multi display scenarios that use different forms of collaborative navigation or environment. If equal partner or more relaxed guided navigation was used, the user may be able to integrate the aid more effectively into their navigational strategies, without having to also use it immediately as part of understanding other's movements. Additional views may still be useful in master/ slave navigation if the environment contained more clearly visible landmarks that exist either naturally within the environment (e.g., within city navigation) or are artificially added.

7.1.3 Contribution 3 - Reduction of breakdowns through path visualisations

Chapter 6 presents the design, technical evaluation and user evaluation of path visualisations, another navigation aid that has previously been successful in desktop navigation. Metrics of path coverage and adjacency were used to evaluate different sampling frequencies of time and spatial resolution (see section 6.2). The application of two types of path visualisation, string and heatmap, to this data was filtered to improve visibility and reduce clutter (see 6.3).

Experiment 4 hypothesised that these path visualisations should reduce the frequency of failure (H4.1) and the provision of a clear path between targets with the string visualisation should reduce the time required to complete the task (H4.2). H4.1 was validated by both path visualisation being beneficial in reducing the frequency of breakdowns (see figure 6.14), and reducing the time taken for participants to complete the task (see figure 6.15).

In contrast, H4.2 was shown to be false, with participants completing the task more quickly in the heatmap condition than in the string condition. One reason for this was that it provided participants with more potential paths to follow (see figure 6.19). Further analysis showed that the visualisations reduced instances of participants getting lost (see figure 6.16), uncertainty surrounding the target area (see figure 6.17) and were closely followed by participants (see figure 6.18).

The effectiveness of these aids indicate that they should be used over additional view aids in similar conditions to this scenario and can be extended to aid longer term collaboration. As this scenario provides a close to worse case scenario for spatial breakdown, it is unlikely that changes in the collaborative environment or dataset will produce conditions that result in this aid being ineffective. It is important to note that these forms of aid do modify the environment and it is unclear what effect these differences will have to the level of content breakdowns present as these were not considered as part of this thesis.

As discussed within the original scenario design this scenario could see frequent and longer term collaboration which can be considered through simple modifications of the data encoded within the path visualisations using the classification within section 2.4.2. This allows encoding of data such as the tracking an individual's membership in different groups and a history of their interactions with the environment. However, this is likely to require more advanced forms of path filtering to ensure that trail pollution does not become an issue.

7.1.4 Contribution 4 - Successful application of heatmaps in aiding navigation

As discussed within the last contribution, participants in Experiment 4 using the heatmap visualisations were shown to be more successful than the base condition (see figure 6.14), complete the task in less time than the base or string conditions (see figure 6.15) and used alternative paths to the string condition (see figure 6.19). These results suggest that navigation aids based on heatmap visualisations should be more frequently considered in conditions that string based aids have been previously successful.

Similarly to the use of general path visualisations, this aid is likely to be successful in scenarios with different collaborative environments or datasets. However, the visual impact of this form of aid on the environment is much greater than with string based visualisations, which may limit some forms of datasets that it can be used with.

7.1.5 Other Findings

Alongside the main findings from the experiments presented above, there are a couple of other findings that have come from this research and are described here.

The first of these is the technical evaluation and filtering methods for path visualizations presented within chapter 6. Previous research into path visualisation has frequently addressed the need to filter the underlying data to achieve an effective visualisation [56, 98], but has rarely considered how differences in the sampling of the raw data affects the resultant visualisations.

The technical evaluation presented in section 6.2 allowed some factors (e.g., position vs. view data) to be evaluated that would otherwise have to be investigated as part of an extensive user study. It also provided empirical evidence for suitable settings of other parameters (e.g., sampling frequency and grid size) which would otherwise have been arbitrarily set and/or developed based on subjective user feedback. The metrics demonstrated here could be used to compare the characteristics of these paths with those of other studies, potentially providing empirical evidence to consider when behaviours such as those discovered through this thesis are likely to occur.

The filtering methods presented within section 6.3 are good examples of methods that have been developed in response to particular problems within this scenario, but could have more universal utility. Non-linear mapping of colour is not novel in itself [13, 108] but provides a simple way to customise the visualisation here to a user's expectations of being able to recognise an area that has been previously visited. The most popular path algorithm addresses 'path pollution' through adapting the commonly used shortest path algorithm to the data that is available.

Secondly, it is worth noting the combination of data, software and hardware that provide the experimental platform as presented in chapter 3. This platform provides an example of how such a platform can be achieved both for further experiments, or more general purposes where additional functionality such as multiple simultaneous users or further path logging could be easily implemented.

The recent pace of development in the area of multiple display arrays has been very rapid, such that elements that previously required specialised software or libraries, can now be achieved using more simple and standardised solutions. For example, the wall-sized display used within this thesis (see section 3.4) could instead be constructed using $\approx$30-inch displays, each with a resolution of 3840 x 2160 pixels arranged three displays wide and two displays tall. This array could be driven through one machine utilising a AMD Firepro W9100 graphics card without the need for synchronised data between multiple machines or specialised libraries such as VRjuggler.

However, more advanced solutions are still required when larger arrays are required such as a fully enclosed room [91]. The design of the current setup can be more easily scaled up to allow such arrays to be implemented.

7.2 Future Work

There are two major forms of future work that could come from the work presented in this thesis, further investigation of navigational aids, as well as implementing the aids that have been developed into wider collaborative systems.

7.2.1 Navigation Aids

The varying effectiveness of the navigation aids previously used in desktop navigation shows that other desktop navigational aids may also be suitable to use in wall-sized to desktop collaboration. However, further experiments would have to validate their effect on performance and behaviour before they could be used. Some of these aids may help address the problems that desktop users had in using some of the aids explored within this research. For example, the difficulties that desktop users had in attending to additional views (see section 4.2.3) could be reduced through the use of additional or artificial landmarks [92] that are clearly visible within those views.

The navigational aids within this research were evaluated primarily with short term tasks which the participants conducted over a maximum of a few hours. Although care was taken to consider some of the issues with longer term use such as filtering the paths used as part of the string path visualisation, the results indicated that participants were willing to use the raw data to increase the efficiency with which they could complete their task (see figure 6.19). With larger quantities of data, trail pollution [97] would reduce the potential to use the raw data in this way and subsequently decrease the effectiveness of those aids, potentially re-introducing breakdowns to the collaboration.

Therefore the effect of the current aids on longer navigational tasks must be evaluated, with potential adaptations (e.g., different forms of automatic trail filtering [98] or fading trails over time [56]) considered to ensure their effectiveness over a longer period of time. Over longer time periods it is also more likely that asynchronous interaction will occur with some of the collaborators wanting to revisit specific areas on their own, a scenario where only a subset of paths may want to be used according to the semantics of the data viewed during the paths [126].

7.2.2 Wider Collaborative Systems and Scenarios

The aids that have been developed within this thesis have only been investigated and evaluated on a subset of navigational tasks which are important to allow users to conduct a wider range of collaborative processes. Further research could investigate whether these aids are sufficient to allow effective collaborative activity to take place between wall-sized and desktop displays and how the use of a wall-sized display benefits the collaborative process. This investigation would also highlight any further collaborative processes that require assistance.

This could be done through extending the scenario examined in this research and conducting a collaborative analysis of a student field trip area before and after visiting the real location. Other scenarios that could be adapted are the areas of co-located sensemaking with wall-sized and desktop/ laptop displays [124] or remote collaboration in immersive systems [67].

Other heterogeneous display scenarios could utilise the findings of this work. For example in a collaboration between desktop and tablet displays there is a smaller difference in display size and FOV than between a wall-sized and desktop display. The use of path visualisations alone may be sufficient to reduce the frequency of common ground breakdowns sufficiently to allow for a working scenario, despite the differences between desktop and tablet input methods.

Appendix A

Participant Information sheet

The following information sheets were given to participants (with appropriate modifications according to the specific experiment).

Participant Information Sheet

Research Project: Collaborative Interaction with Heterogeneous Displays Experiment 4: Powerwall/Desktop Ball of String vs Heatmap path visualisations

This sheet will hopefully provide you with enough information about the study to allow you to make an informed decision about participation. However, if you have any questions or would like to discuss anything with me please let me know.

What is the project's purpose?

I am a PhD Student in the School of Computing at the University of Leeds. The aim of my research is to develop interaction techniques that help people perform collaborative work with Powerwall (large, high resolution) and standard desktop displays.

Do I have to take part?

This research is subject to ethical guidelines set out by the University. These guidelines include principles such as obtaining your informed consent before research starts, notifying you of your right to withdraw at any time, and protection of your anonymity. You should not take part in this study if you suffer from:

- Epilepsy
- Migraines

What will happen to me if I take part?

The experiment will involve 3 sessions, each of approximately one hour, during which rests will be provided to prevent fatigue. The experiment takes place in a virtual 3D landscape. A participant's task in each trial is to observe while they are taken to "target" places and then, having been put back at the start point, navigate themselves to the targets.

The experiment procedure for each session is as follows:

- Experimenter will demonstrate the user interface
- You will practice navigating around the 3D landscape
- Perform 8 practice trials. A trial will be terminated if you have not found the targets after 3 minutes.
- Rest (5 minutes)
- Perform 16 test trials.

Will I be recorded, and how will the recorded media be used?

Your time, end-point accuracy, and the movements you make in the 3D landscape will be recorded for subsequent analysis.

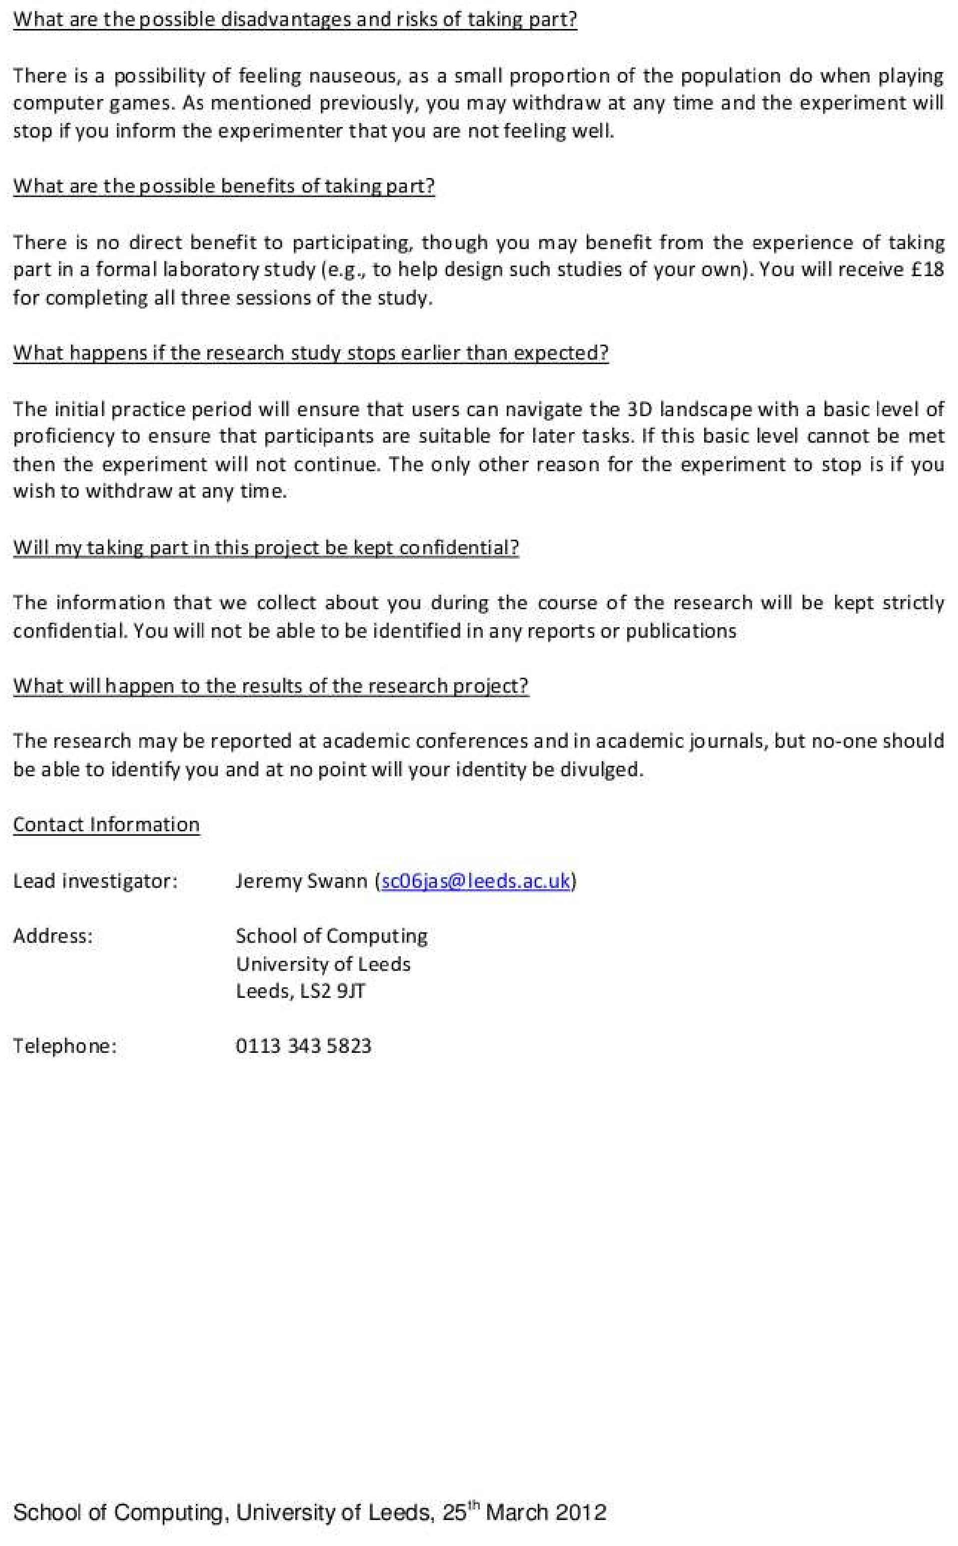

What are the possible disadvantages and risks of taking part?

There is a possibility of feeling nauseous, as a small proportion of the population do when playing computer games. As mentioned previously, you may withdraw at any time and the experiment will stop if you inform the experimenter that you are not feeling well.

What are the possible benefits of taking part?

There is no direct benefit to participating, though you may benefit from the experience of taking part in a formal laboratory study (e.g., to help design such studies of your own). You will receive £18 for completing all three sessions of the study.

What happens if the research study stops earlier than expected?

The initial practice period will ensure that users can navigate the 3D landscape with a basic level of proficiency to ensure that participants are suitable for later tasks. If this basic level cannot be met then the experiment will not continue. The only other reason for the experiment to stop is if you wish to withdraw at any time.

Will my taking part in this project be kept confidential?

The information that we collect about you during the course of the research will be kept strictly confidential. You will not be able to be identified in any reports or publications

What will happen to the results of the research project?

The research may be reported at academic conferences and in academic journals, but no-one should be able to identify you and at no point will your identity be divulged.

Contact Information

Lead investigator:	Jeremy Swann (sc06jas@leeds.ac.uk)
Address:	School of Computing University of Leeds Leeds, LS2 9JT
Telephone:	0113 343 5823

School of Computing, University of Leeds, 25th March 2012

Appendix B

Participant Consent Form

Participants were asked to read and sign the following form before the experiment began.

University of Leeds

Participant Consent Form

Title of Research Project: Collaborative Interaction with Heterogeneous Displays

Name of Researcher: Mr Jeremy Swann

Initial the box if you agree with the statement to the left

1 I understand that I am free to choose not to answer a question on the consent form without having to give a reason why. ☐

2 I confirm that I have read and understand the information sheet dated 1 September 2011 explaining the above research project and I have had the opportunity to ask questions about the project. ☐

3 I do not suffer from epilepsy and/or migraines. ☐

4 I understand that my participation is voluntary and that I am free to withdraw at any time without giving any reason and without there being any negative consequences. In addition, should I not wish to answer any particular question or questions, I am free to decline. ☐

5 I understand that my responses will be kept strictly confidential. I give permission for members of the research team to have access to my anonymised responses. I understand that my name will not be linked with the research materials, and I will not be identified or identifiable in the report or reports that result from the research. ☐

6 I agree for the data collected from me to be used in future research ☐

7 I agree to take part in the above research project. ☐

________________________	________________	____________________
Name of participant (*or legal representative*)	Date	Signature
________________________	________________	____________________
Name of person taking consent (*if different from lead researcher*) *To be signed and dated in presence of the participant*	Date	Signature
________________________	________________	____________________
Lead researcher *To be signed and dated in presence of the participant*	Date	Signature

Date: ____________ Name of Applicant: ________________________

Appendix C

Navigational Control Sheets

The following information sheets were given to participants using the heatmap visualisation in experiment 4 to explain the controls. Similar sheets were given depending on the specifics of the experiment run.

Navigational Controls

Move

↔ - Left/ right
↕ - forwards/ backwards

Progress through trials

Rotate

↔ - heading
↕ - pitch

Move

↕ - Ascend/ Descend

For each trial pressing button 2 will:

1. Go to start point
2. Start movement
3. When movement is finished – return to start point and begin trial
4. Confirm a target
5. When targets are found or time expired – clear screen (return to 1)

Interface

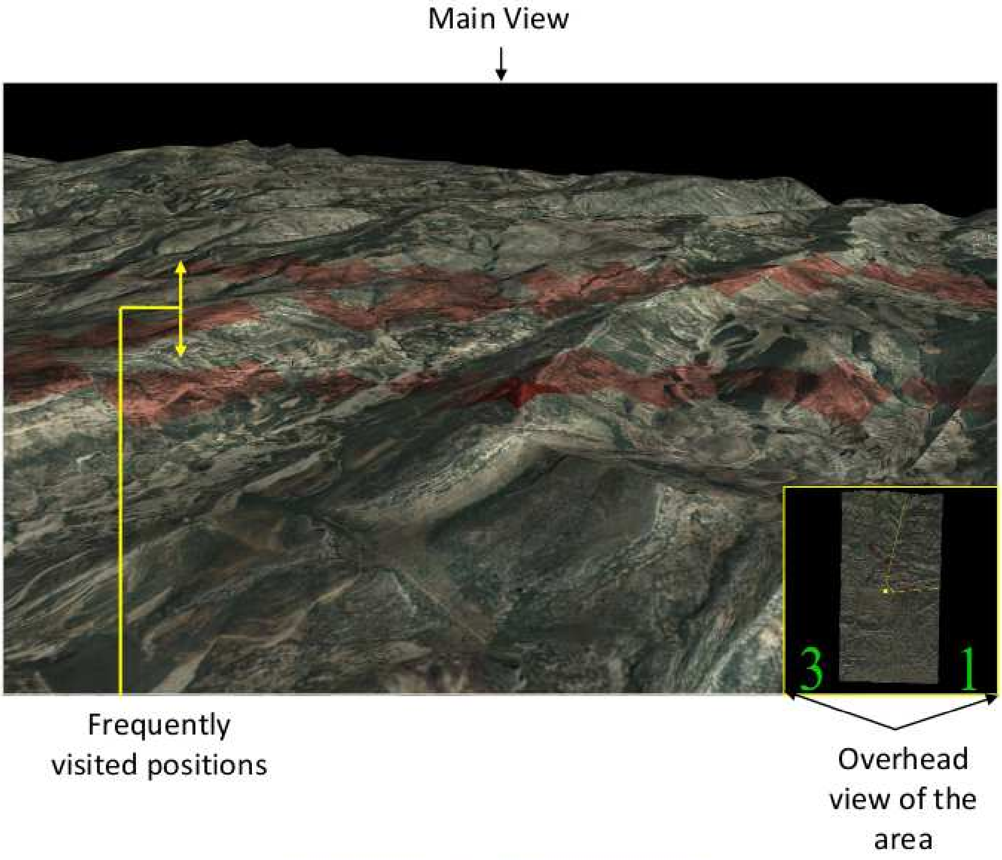
Main View
3
1
Frequently visited positions
Overhead view of the area


D

References

[1] J. Alexander, A. Cockburn, S. Fitchett, C. Gutwin, and S. Greenberg. Revisiting Read Wear: Analysis, Design, and Evaluation of a Footprints Scrollbar. In *CHI '09: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1665–1674, 2009.

[2] E. J. H. Andriessen. *Working with Groupware: understanding and evaluating collaboration technology*. Springer-Verlag, London, England, 1st edition, 2003.

[3] A. J. Aretz. The design of electronic map displays. *Human Factors*, 33(1):pages 85–101, 1991.

[4] M. Argyle and J. Graham. The central Europe experiment: Looking at persons and looking at objects. 1(1):6–16, 1976.

[5] Ascension. Ascension technology corporation. `http://www.ascension-tech.com/realtime/rtflockofbirds.php`. [Online; accessed 14 July 2013].

[6] J.M. Atkinson and J. Heritage. *Structures of social action: Studies in conversation analysis*. Cambridge Univ Press, Cambridge, England, 1984.

[7] A. S. Axelsson. *The digital divide: status differences in virtual environments*, pages 188–204. Springer-Verlag, New York, NY, 2002.

[8] A. Baddeley, S. D. Sala, and T. W. Robbins. Working Memory and Executive Control. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 351(1346):pages 1397–1404, 1996.

[9] R. Ball, C. North, and D. A. Bowman. Move to improve: promoting physical navigation to increase user performance with large displays. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 191–200, 2007.

[10] A. Bangerter. Using Pointing and Describing to Achieve Joint Focus of Attention in Dialogue. *Psychological Science*, 15(6):415–419, June 2004.

[11] J. Bardram. Designing for the dynamics of cooperative work activities. In *CSCW '98: Proceedings of the ACM conference on Computer supported cooperative work*, pages 89–98, 1998.

[12] A. T. Bennett. Do animals have cognitive maps? *The Journal of experimental biology*, 199(1):219–224, 1996.

[13] L. D. Bergman, B. Rogowitz, and L. A. Treinish. A rule-based tool for assisting colormap selection. In *VIS '95: Proceedings of the IEEE conference on Visualization '05*, pages 118–125, 1995.

[14] X. Bi and R. Balakrishnan. Comparing usage of a large high-resolution display to single or dual desktop displays for daily work. In *CHI '09: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1005–1014, 2009.

[15] A. Bierbaum, C. Just, P. Hartling, K. Meinert, A. Baker, and C. Cruz-Neira. VR Juggler: a virtual platform for virtual reality application development. In *VR '01: Proceedings of the IEEE conference on Virtual Reality*, pages 89–96, 2001.

[16] S. A. Bly, S. R. Harrison, and S. Irwin. Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment. *Communications of the ACM*, 36(1):28–46, 1993.

[17] U. M. Borghoff and J. H. Schlichter. *Computer-Supported Cooperative Work: Introduction to Distributed Applications*. Springer, 1st edition, 2000.

[18] K. Börner and S. Penumarthy. Social diffusion patterns in three-dimensional virtual worlds. *Information Visualization*, 2(3):pages 182–198, 2003.

[19] D. A. Bowman, E. T. Davis, L. F. Hodges, and A. N. Badre. Maintaining Spatial Orientation during Travel in an Immersive Virtual Environment. *Presence: Teleoperators and Virtual Environments*, 8(6):pages 618–631, 1999.

[20] D. A. Bowman, E. Kruijff, J. J. LaViola, and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, Boston, MA, 1st edition, 2004.

[21] R. Brown, S. Joslin, and P. Drennan. A Visualisation Framework for Collaborative Virtual Environment Usage Information. In R. Liere, T. Adriaansen, and E. Zudilova-Seinstra, editors, *Trends in Interactive Visualization*, Advanced Information and Knowledge Processing, pages 201–221. Springer, London, England, 2009.

[22] N. Burtnyk, A. Khan, G. Fitzmaurice, R. Balakrishnan, and G. Kurtenbach. StyleCam: interactive stylized 3D navigation using integrated spatial & temporal controls. In *UIST '02: Proceedings of the ACM symposium on User interface software and technology*, pages 101–110, New York, NY, USA, 2002.

[23] J. M. Carroll, D. C. Neale, P. L. Isenhour, M. B. Rosson, and D. S. McCrickard. Notification and awareness: synchronizing task-oriented collaborative activity. *International Journal of Human-Computer Studies*, 58(5):605–632, 2003.

[24] M. Chalmers, R. Ingram, and C. Pfranger. Adding Imageability Features to Information Displays. In *UIST '96: Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 33–39, 1996.

[25] L. Chittaro, R. Ranon, and L. Ieronutti. VU-Flow: A Visualization Tool for Analyzing Navigation in Virtual Environments. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):pages 1475–1485, 2006.

[26] E. Chrastil and W. Warren. Testing models of path integration in a triangle completion task. *Journal of Vision*, 8:1153, 2008.

[27] M. C. Chuah and S. F. Roth. Visualizing Common Ground. In *IV '03: Proceedings of the IEEE International Conference on Information Visualization*, pages 365–372, 2003.

[28] H. H. Clark and S. A. Brennan. Grounding in communication. In L. B. Resnick, J. M. Levine, and S. D. Teasley, editors, *Perspectives on socially shared cognition*, pages 127–149. APA Books, Washington, WA, 1991.

[29] H. H. Clark and D. Wilkes-Gibbs. Referring as a collaborative process. *Cognition*, 22(1):1–39, 1986.

[30] G. Convertino, C. H. Ganoe, W. A. Schafer, B. Yost, and J. M. Carroll. A Multiple View Approach to Support Common Ground in Distributed and Synchronous Geo-Collaboration. In *CMV '05: Proceedings of the IEEE conference on Coordinated and Multiple Views in Exploratory Visualization*, pages 121–132. IEEE, 2005.

[31] G. Convertino, H. M. Mentis, M. B. Rosson, J. M. Carroll, A. Slavkovic, and C. H. Ganoe. Articulating Common Ground in Cooperative Work: Content and Process. In *CHI '08: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1637–1646, 2008.

[32] G. Convertino, H. M. Mentis, M. B. Rosson, A. Slavkovic, and J. M. Carroll. Supporting Content and Process Common Ground in Computer-supported Teamwork. In *CHI '09: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2339–2348, 2009.

[33] G. Convertino, H. M. Mentis, A. Slavkovic, M. B. Rosson, and J. M. Carroll. Supporting common ground and awareness in emergency management planning: A design research project. *ACM Transactions on Computer-Human Interaction*, 18(4), 2011.

[34] M. Czerwinski, D. S. Tan, and G. G. Robertson. Women take a wider view. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 195–202, 2002.

[35] R. P. Darken and B. Peterson. Spatial Orientation, Wayfinding, and Representation. In K. Stanney, editor, *Handbook of Virtual Environment Technology*, pages 493–518. Lawrence Erlbaum Associates, Mahwah, NJ, 1st edition, 2002.

[36] R. P. Darken and J. L. Sibert. A toolset for navigation in virtual environments. In *UIST '93: Proceedings of the 6th annual ACM symposium on User interface software and technology*, pages 157–165, 1993.

[37] R. P. Darken and J. L. Sibert. Navigating large virtual spaces. *International Journal of Human-Computer Interaction*, 8(1):49–71, 1996.

[38] S. de Groot, F. Centeno Ricote, and J. C. F. de Winter. The effect of tire grip on learning driving skill and driving style: A driving simulator study. *Transportation Research Part F: Traffic Psychology and Behaviour*, 15(4):413–426, 2012.

[39] D. DeRoure, W. Hall, S. Reich, G. Hill, A. Pikrakis, and M. Stairmand. MEMOIR an open framework for enhanced navigation of distributed information. *Information Processing and Management*, 37(1):pages 53–74, 2001.

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

[41] A. Dix, J. E. Finlay, G. D. Abowd, and R. Beale. *Human-Computer Interaction*. Prentice Hall, Harlow, England, 3rd edition, 2003.

[42] T. J. Dodds and R. A. Ruddle. Mobile group dynamics in Large-Scale collaborative virtual environments. In *VR '08: Proceedings of the IEEE Virtual Reality conference*, pages 59–66, 2008.

[43] T. J. Dodds and R. A. Ruddle. Using teleporting, awareness and multiple views to improve teamwork in collaborative virtual environments. In *EGVE '08: Proceedings of the Eurographics symposium on virtual environments*, pages 81–8, 2008.

[44] T. J. Dodds and R. A. Ruddle. Using mobile group dynamics and virtual time to improve teamwork in large-scale collaborative virtual environments. *Computers & Graphics*, 33(2):pages 130–138, 2009.

[45] K. Doerr and F. Kuester. CGLX: A Scalable, High-Performance Visualization Framework for Networked Display Environments. *IEEE Transactions on Visualization and Computer Graphics*, 17(3):pages 320–332, 2011.

[46] G. Drettakis, M. Roussou, A. Reche, and N. Tsingos. Design and Evaluation of a Real-World Virtual Environment for Architecture and Urban Planning. *Presence: Teleoperators and Virtual Environments*, 16(3):318–332, 2007.

[47] J. Dyck and C. Gutwin. Groupspace: a 3D workspace supporting user awareness. In *CHI EA '02: Extended abstracts of the SIGCHI conference on Human factors in computing systems*, pages 502–503, 2002.

[48] N. Elmqvist, M. E. Tudoreanu, and P. Tsigas. Evaluating motion constraints for 3D wayfinding in immersive and desktop virtual environments. In *CHI '08: Proceeding of the SIGCHI conference on Human factors in computing systems*, pages 1769–1778, 2008.

[49] T. T. Elvins, D. R. Nadeau, R. Schul, and D. Kirsh. Worldlets: 3-D Thumbnails for Wayfinding in Large Virtual Worlds. *Presence: Teleoperators and Virtual Environments*, 10(6):pages 565–582, 2001.

[50] Blizzard Entertainment. World of Warcraft, 2004.

[51] D. Fisher. Hotmap: Looking at Geographic Attention. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1184–1191, 2007.

[52] P. Foo, W. H. Warren, A. Duchon, and M. J. Tarr. Do humans integrate routes into a cognitive map? Map- versus landmark-based navigation of novel shortcuts. *Journal of experimental psychology. Learning, memory, and cognition*, 31(2):195–215, 2005.

[53] N. Foreman, D. Foreman, A. Cummings, and S. Owens. Locomotion, Active Choice, and Spatial Memory in Children. *The Journal of General Psychology*, 117(2):215–235, 1990.

[54] A. Ginige, M. Romano, M. Sebillo, G. Vitiello, and P. Di Giovanni. Spatial data and mobile applications: general solutions for interface design. In *AVI '12: Proceedings of the ACM International Working Conference on Advanced Visual Interfaces*, pages 189–196, 2012.

[55] C. Goodyer, J. Hodrien, J. Wood, P. Kohl, and K. Brodlie. Using high-resolution displays for high-resolution cardiac data. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1898):pages 2667–2677, 2009.

[56] D. Grammenos, A. Mourouzis, and C. Stephanidis. Virtual prints: Augmenting virtual environments with interactive personal marks. *International Journal of Human-Computer Studies*, 64(3):pages 221–239, 2006.

[57] S. Greenberg, C. Gutwin, and M. Roseman. Semantic telepointers for groupware. In *Proceedings of the IEEE Australian Conference on Computer-Human Interaction*, pages 54–61, 1996.

[58] H. P Grice. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics Volume 3: Speech Acts*, pages 225–242. Academic Press, New York, NY, 1st edition, 1975.

[59] T. Gross. Supporting Effortless Coordination: 25 Years of Awareness Research. 22(4-6):425–474, 2013.

[60] J. Grudin. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In *CHI '01: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 458–465, 2001.

[61] C. Gutwin and S. Greenberg. Workspace Awareness for Groupware. In *CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 208–209, 1996.

[62] C. Gutwin and S. Greenberg. The Effects of Workspace Awareness Support on the Usability of Real-time Distributed Groupware. *ACM Transactions on Computer-Human Interaction*, 6(3):243–281, 1999.

[63] C. Gutwin and S. Greenberg. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work (CSCW)*, 11(3):411–446, 2002.

[64] C. Gutwin, S. Greenberg, and M. Roseman. Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. In M. A. Sasse, J. Cunningham, and R. L. Winder, editors, *People and Computers XI*, pages 281–298. Springer London, 1996.

[65] C. Gutwin and R. Penner. Improving Interpretation of Remote Gestures with Telepointer Traces. In *CSCW '02: Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pages 49–57, 2002.

[66] M. Hartswood and R. Procter. Design guidelines for dealing with breakdowns and repairs in collaborative work settings. *International Journal of Human-Computer Studies*, 53(1):pages 91–120, 2000.

[67] I. Heldal, A. Steed, M. Spante, R. Schroeder, S. Bengtsson, and M. Partanen. Successes and failures in Co-Present situations. *Presence: Teleoperators and Virtual Environments*, 14(5):pages 563–579, 2005.

[68] W. C. Hill, J. D. Hollan, D. Wroblewski, and T. McCandless. Edit Wear and Read Wear. In *CHI '92: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3–9, 1992.

[69] J. Hindmarsh, M. Fraser, C. Heath, S. Benford, and C. Greenhalgh. Object-focused interaction in collaborative virtual environments. *ACM Transactions on Computer-Human Interaction*, 7(4):pages 477–509, 2000.

[70] N. Hoobler, G. Humphreys, and M. Agrawala. Visualizing Competitive Behaviors in Multi-User Virtual Environments. In *VIS '04: Proceedings of the IEEE conference on Visualization '04*, pages 163–170, 2004.

[71] D. Iaboni and C. MacGregor. The Design of a Virtual Trailblazing Tool Virtual and Mixed Reality. In Randall Shumaker, editor, *Virtual and Mixed Reality*, volume 5622 of *Lecture Notes in Computer Science*, chapter 22, pages 186–195. Springer Berlin / Heidelberg, 2009.

[72] J. Jacobson and M. Lewis. An Experimental Comparison Of Three Methods For Collision Handling In Virtual Environments. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, pages 1273–1277, 1997.

[73] B. Jeong, L. Renambot, R. Jagodic, R. Singh, J. Aguilera, A. Johnson, and J. Leigh. High-Performance Dynamic Graphics Streaming for Scalable Adaptive Graphics Environment. In *SC '06: Proceedings of the ACM/IEEE Supercomputing Conference*, page 24, 2006.

[74] R. Johansen. *GroupWare: Computer Support for Business Teams*. The Free Press, New York, NY, 1st edition, 1988.

[75] F. Kappe, B. Zaka, and M. Steurer. Automatically Detecting Points of Interest and Social Networks from Tracking Positions of Avatars in a Virtual World. In *Proceedings of the IEEE International Conference on Advances in Social Network Analysis and Mining*, pages 89–94, 2009.

[76] S. Lessels and R. A. Ruddle. Changes in navigational behaviour produced by a wide field of view and a high fidelity visual scene. In *EGVE '04: Proceedings of the 10th Eurographics Symposium on Virtual Environments*, pages 71–78, 2004.

[77] J. Li, S. Greenberg, E. Sharlin, and J. Jorge. Interactive Two-sided Transparent Displays: Designing for Collaboration. In *DIS '14: Proceedings of the Conference on Designing Interactive Systems*, pages 395–404, 2014.

[78] J. M. Loomis, R. L. Klatzky, R. G. Golledge, and J. W. Philbeck. Human navigation by path integration. In R. Golledge, editor, *Wayfinding: Cognitive mapping and other spatial processes*, pages 125–151. Johns Hopkins University Press, Baltimore, MD, 1999.

[79] K. Lynch. *The Image of the City*. The MIT Press, Cambridge, MA, 1st edition, 1960.

[80] Microsoft. Xbox Kinect Motion Sensors - Games, Photos & News — Xbox.com UK - Xbox.com. `http://www.xbox.com/en-GB/Kinect`. [Online; accessed 10 July 2013].

[81] B. A. Nardi, editor. *Context and Consciousness: Activity Theory and Human-Computer Interaction*. The MIT Press, Cambridge, MA, 1995.

[82] D. C. Neale, J. M. Carroll, and M. B. Rosson. Evaluating computer-supported cooperative work: models and frameworks. In *CSCW '04: Proceedings of the ACM conference on Computer Supported Cooperative Work*, pages 112–121, 2004.

[83] C. C. Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The CAVE: audio visual experience automatic virtual environment. *Communications of the ACM*, 35(6):pages 64–72, 1992.

[84] G. M. Olson and J. S. Olson. Distance matters. *Human- Computer Interaction*, 15(2):pages 139–178, 2000.

[85] A. Parush and D. Berman. Navigation and orientation in 3D user interfaces: the impact of navigation aids and landmarks. *International Journal of Human-Computer Studies*, 61(3):375–395, 2004.

[86] L. Pearce, P. Route maps: A study of travellers' perceptions of a section of countryside. *Journal of Environmental Psychology*, 1(2):141–155, 1981.

[87] E. R. Pedersen and Tomas S. AROMA: Abstract Representation of Presence Supporting Mutual Awareness. In *CHI '97: Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pages 51–58, 1997.

[88] V. Penichet, I. Marin, J. Gallud, M. Lozano, and R. Tesoriero. A Classification Method for CSCW Systems. In *VODCA '06: Proceedings of the International Workshop on Views on Designing Complex Architectures*, pages pages 237–247. 2007.

[89] J. Peponis, R. C. Dalton, J. Wineman, and N. Dalton. Measuring the effects of layout upon visitors' spatial behaviors in open plan exhibition settings. *Environment and Planning B: Planning and Design*, 31(3):pages 453–473, 2004.

[90] P. Peruch, J. Vercher, and G. M. Gauthier. Acquisition of Spatial Knowledge Through Visual Exploration of Simulated Environments. *Ecological Psychology*, 7(1):1–20, 1995.

[91] K. Petkov, C. Papadopoulos, and A. E. Kaufman. Visual exploration of the infinite canvas. In *VR '13: Proceedings of the IEEE Virtual Reality conference*, pages 11–14, 2013.

[92] J. S. Pierce and R. Pausch. Navigation with place representations and visible landmarks. In *VR '04: Proceedings of the IEEE Conference on Virtual Reality*, pages 173–288, 2004.

[93] A. E. Richardson, D. R. Montello, and M. Hegarty. Spatial knowledge acquisition from maps and from navigation in real and virtual environments. *Memory & Cognition*, 27(4):pages 741–750, 1999.

[94] D. Roberts, R. Wolff, O. Otto, and A. Steed. Constructing a Gazebo: Supporting Teamwork in a Tightly Coupled, Distributed Task in Virtual Reality. *Presence: Teleoperators and Virtual Environments*, 12(6):644–657, 2003.

[95] M. Rohs, J. Schöning, M. Raubal, G. Essl, and A. Krüger. Map navigation with mobile devices: virtual versus physical movement with and without visual context. In *ICMI '07: Proceedings of the ACM international conference on Multimodal interfaces*, pages 146–153, 2007.

[96] C. Rooney and R. A. Ruddle. Improving Window Manipulation and Content Interaction on High-Resolution, Wall-Sized Displays. *International Journal of Human-Computer Interaction*, 28(7):pages 423–432, 2011.

[97] R. A. Ruddle. The Effect of Trails on First-time and Subsequent Navigation in a Virtual Environment. In *VR '05: Proceedings of the IEEE Conference on Virtual Reality*, pages 115 – 122, 2005.

[98] R. A. Ruddle. Generating Trails Automatically, to Aid Navigation When You Revisit an Environment. *Presence: Teleoperators and Virtual Environments*, 17(6):pages 562–574, 2008.

[99] R. A. Ruddle and D. M. Jones. Movement in Cluttered Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 10(5):511–524, 2001.

[100] R. A. Ruddle and S. Lessels. Three Levels of Metric for Evaluating Wayfinding. *Presence: Teleoperators and Virtual Environments*, 15(6):637–654, 2006.

[101] R. A. Ruddle, S. J. Payne, and D. M. Jones. Navigating Large-Scale "Desk-Top" Virtual Buildings: Effects of Orientation Aids and Familiarity. *Presence: Teleopertors and Virtual Environments*, 7:pages 179–192, 1998.

[102] R. A. Ruddle, S. J. Payne, and D. M. Jones. The Effects of Maps on Navigation and Search Strategies in Very-Large-Scale Virtual Environments. *Journal of Experimental Psychology: Applied*, 5(1):pages 54–75, 1999.

[103] R. A. Ruddle and P. Péruch. Effects of proprioceptive feedback and environmental characteristics on spatial learning in virtual environments. *International Journal of Human-Computer Studies*, 60(3):299–326, 2004.

[104] R. A. Ruddle, J. C. D. Savage, and D. M. Jones. Verbal communication during cooperative object manipulation. In *CVE '02: Proceedings of the ACM international conference on Collaborative virtual environments*, CVE '02, pages 120–127, 2002.

[105] R. A. Ruddle, J. C. D. Savage, and D. M. Jones. Levels of Control During a Collaborative Carrying Task. *Presence: Teleoperators and Virtual Environments*, 12(2):pages 140–155, 2003.

[106] R. A. Ruddle, E. Volkova, and H. H. Bülthoff. Walking improves your cognitive map in environments that are large-scale and large in extent. *ACM Transactions on Computer-Human Interaction*, 18, 2011.

[107] R. A. Ruddle, E. Volkova, B. Mohler, and H. H. Bülthoff. The effect of landmark and body-based sensory information on route knowledge. *Memory & Cognition*, 39(4):pages 686–699, 2011.

[108] P. Schulze-Wollgast, C. Tominski, and H. Schumann. Enhancing visual exploration by appropriate color coding. In *WSCG '05: Proceedings of the International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*, pages 203–210, 2005.

[109] S. D. Scott, M. Sheelagh, T. Carpendale, and K. M. Inkpen. Territoriality in Collaborative Tabletop Workspaces. In *CSCW '04: Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pages 294–303, 2004.

[110] A. W. Siegel and S. H. White. The Development of Spatial Representations of Large-Scale Environments. volume 10 of *Advances in Child Development and Behaviour*, pages 9–55. Elsevier, 1975.

[111] M. Slater, A. Sadagic, M. Usoh, and R. Schroeder. Small-Group behavior in a virtual and real environment: A comparative study. *Presence: Teleoperators and Virtual Environments*, 9(1):pages 37–51, 2000.

[112] M. Spante, R. Schroeder, A.S. Axelsson, and M. Christie. How putting yourself into the other persons virtual shoes enhances collaboration. In *PRESCENCE '04: Proceedings of the International Workshop on Presence*, pages 190–196, 2004.

[113] A. P. Steed, D. J. Roberts, R. Schroeder, and I. Heldal. Interaction between users of immersion projection technology systems. In *HCI '11: Proceedings of the International Conference on Human Computer Interaction (HCI)*, 2005.

[114] M. Stefik, D. G. Bobrow, G. Foster, S. Lanning, and D. Tatar. WYSIWIS revised: early experiences with multiuser interfaces. *ACM Transactions on Information Systems*, 5(2):147–167, 1987.

[115] C. Steinfield, C. Y. Jang, and B. Pfaff. Supporting virtual team collaboration: the TeamSCOPE system. In *GROUP '99: Proceedings of the international ACM SIGGROUP conference on Supporting group work*, pages 81–90, 1999.

[116] I. Steinmacher, A. P. Chaves, and M. G. Gerosa. Awareness Support in Distributed Software Development: A Systematic Review and Mapping of the Literature. *Computer Supported Cooperative Work (CSCW)*, 22(2-3):113–158, 2013.

[117] E. A. Suma, S. L. Finkelstein, M. Reid, S. V. Babu, A. C. Ulinski, and L. F. Hodges. Evaluation of the Cognitive Effects of Travel Technique in Complex

Real and Virtual Environments. *IEEE Transactions on Visualization and Computer Graphics*, 16(4):pages 690–702, 2010.

[118] D. S. Tan, D. Gergle, P. G. Scupelli, and R. Pausch. Physically Large Displays Improve Path Integration in 3D Virtual Navigation Tasks. In *CHI '04: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 439–446, 2004.

[119] Unity Technologies. Unity - Game Engine. `http://unity3d.com/`. [Online; accessed 15 July 2013].

[120] E. Teh, S. Jamson, O. Carsten, and H. Jamson. Temporal fluctuations in driving demand: The effect of traffic complexity on subjective measures of workload and driving performance. *Transportation Research Part F: Traffic Psychology and Behaviour*, 22:207–217, 2014.

[121] P. W. Thorndyke and B. Hayes-Roth. Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14(4):pages 560–589, 1982.

[122] D. Treanor, N. Jordan-Owers, J. Hodrien, J. Wood, P. Quirke, and R. A. Ruddle. Virtual reality Powerwall versus conventional microscope for viewing pathology slides: an experimental comparison. *Histopathology*, 55(3):pages 294–300, 2009.

[123] Midland Valley. software : Move : Midland Valley : The Structural Geology Experts. `http://www.mve.com/software/move`. [Online; accessed 9 Sept 2011].

[124] K. Vogt, L. Bradel, C. Andrews, C. North, A. Endert, and D. Hutchings. Co-located collaborative sensemaking on a large high-resolution display with multiple input devices. In *INTERACT '11: Proceedings of the IFIP international conference on Human-computer interaction*, pages 589–604, 2011.

[125] A. Voida, S. Voida, S. Greenberg, and H. A. He. Asymmetry in media spaces. In *CSCW '08: Proceedings of the 2008 ACM conference on Computer supported cooperative work*, pages 313–322, 2008.

[126] S. Vosinakis and I. Papadakis. Virtual Worlds as Information Spaces: Supporting Semantic and Social Navigation in a shared 3D Environment. In *VS-GAMES '11: Proceedings of the IEEE International Conference on Games and Virtual Worlds for Serious Applications (VS-GAMES)*, pages 220–227, 2011.

[127] Oculus VR. Oculus Rift - Virtual Reality Headset for 3D Gaming — Oculus VR. `http://www.oculusvr.com/`. [Online; accessed 10 July 2013].

[128] W. A. Wagenaar. My memory: A study of autobiographical memory over six years. *Cognitive Psychology*, 18(2):pages 225–252, 1986.

[129] D. Waller. Individual differences in spatial learning from computer-simulated environments. *Journal of experimental psychology. Applied*, 6(4):pages 307–321, 2000.

[130] D. Waller, E. Hunt, and D. Knapp. The Transfer of Spatial Knowledge in Virtual Environment Training. *Presence: Teleoperators and Virtual Environments*, 7(2):129–143, 1998.

[131] D. L. Weatherford. *Representing and manipulating spatial information from different environments: Models to neighborhoods*, pages 41–70. Lawrence Erlbaum Associates, Mahwah, NJ, 1985.

[132] C. D. Wickens and T. T. Prevett. Exploring the Dimensions of Egocentricity in Aircraft Navigation Displays. *Journal of Experimental Psychology: Applied*, 1(2):pages 110–135, 1995.

[133] P. N. Wilson, N. Foreman, R. Gillett, and D. Stanton. Active Versus Passive Processing of Spatial Information in a Computer-Simulated Environment. *Ecological Psychology*, 9(3):207–222, 1997.

[134] N. Wong and C. Gutwin. Where are you pointing?: the accuracy of deictic pointing in CVEs. In *CHI '10: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1029–1038, 2010.

[135] Y. L. Wong, J. Zhao, and N. Elmqvist. Evaluating Social Navigation Visualization in Online Geographic Maps. *International Journal of Human-Computer Interaction*, 31(2):118–127, 2014.

[136] C. A. Zanbaka, B. C. Lok, S. V. Babu, A. C. Ulinski, and L. F. Hodges. Comparison of path visualizations and cognitive measures relative to travel technique in a virtual environment. *IEEE Transactions on Visualization and Computer Graphics*, 11(6):pages 694–705, 2005.