Augmenting User Interfaces with Haptic Feedback

Christopher Trevor Asque

A thesis submitted for the degree of
Doctor of Philosophy
at the University of East Anglia
April 2014

Augmenting User Interfaces with Haptic Feedback

Christopher Trevor Asque

© This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with the author and that no quotation from the thesis, nor any information derived therefrom, may be published without the author's prior, written consent.

Abstract

Computer assistive technologies have developed considerably over the past decades. Advances in computer software and hardware have provided motion-impaired operators with much greater access to computer interfaces. For people with motion impairments, the main difficulty in the communication process is the input of data into the system. For example, the use of a mouse or a keyboard demands a high level of dexterity and accuracy. Traditional input devices are designed for able-bodied users and often do not meet the needs of someone with disabilities. As the key feature of most graphical user interfaces (GUIs) is to point-and-click with a cursor this can make a computer inaccessible for many people.

Human-computer interaction (HCI) is an important area of research that aims to improve communication between humans and machines. Previous studies have identified haptics as a useful method for improving computer access. However, traditional haptic techniques suffer from a number of shortcomings that have hindered their inclusion with real world software. The focus of this thesis is to develop haptic rendering algorithms that will permit motion-impaired operators to use haptic assistance with existing graphical user interfaces. The main goal is to improve interaction by reducing error rates and improving targeting times. A number of novel haptic assistive techniques are presented that utilise the three degrees-of-freedom (3DOF) capabilities of modern haptic devices to produce assistance that is designed specifically for motion-impaired computer users. To evaluate the effectiveness of the new techniques a series of point-and-click experiments were undertaken in parallel with cursor analysis to compare the levels of performance. The task required the operator to produce a predefined sentence on the densely populated Windows on-screen keyboard (OSK). The results of the study prove that higher performance levels can be

achieved using techniques that are less constricting than traditional assistance.

Acknowledgements

Firstly, I would like to thank my Supervisors Dr. Stephen Laycock and Professor Andy Day for their support and guidance throughout the course of the project. Their advice, enthusiasm and academic experience, have been invaluable to me.

I owe my deepest gratitude to the volunteers and staff at NANSA with special thanks to Richard Shepherd and Robert Dent for their advice and suggestions. Without the volunteers' commitment towards performing experimental tasks this thesis would not have been possible.

I would like to acknowledge with tremendous and deep thanks my family and my wife Sarah. Through their love, patience, support and unwavering belief in me, I have been able to complete my academic studies. They have taught me many important skills that have helped me progress in life and my career. I would like to take this opportunity to thank my brother-in-law Dr. Richard Ingram for his advice regarding the statistical analysis. I would also like to thank my good friends Liam Gill and Ash Ellis for providing me with accommodation during the start of my PhD. Special thanks to Liam for his advice and C# expertise in regard to the feature extractor.

Finally, I would like to thank the CMP support team at UEA. The team are a great asset to the computing department and have helped to resolve many software and hardware issues encountered over the course of the project. Special thanks to Peter Trollope, Matthew Ladd, Binoop Pulikkottil and Russell Smith.

Publications

Several parts of this thesis have been published at the following venues:

- Christopher T. Asque, Andy M. Day, Stephen D. Laycock, *"Haptic Assisted Target Acquisition in a Visual Point-and-Click Task for Computer Users with Motion-Impairments"*, IEEE Transactions on Haptics, 05 Oct. 2011. IEEE computer Society Digital Library. IEEE Computer Society.
- Christopher T. Asque, Andy M. Day, Stephen D. Laycock, *"Cursor Navigation using Haptics for Motion-Impaired Computer Users"*, Haptics: Perception, Devices, Mobility and Communication, International Conference, EuroHaptics 2012, Tampere, Finland, June 13-15, 2012, 13-24.
- Christopher T. Asque, Andy M. Day, Stephen D. Laycock, *"Augmenting Graphical User Interfaces with Haptic Assistance for Motion-Impaired Operators"*, To appear in the International Journal of Human-Computer Studies, 2014.

Table of Contents

List of Tables

List of Figures

List of Algorithms

Chapter 1

Introduction

1.1 Background

The functionality of a computer offers many benefits for disabled users. Computer access for people with motion impairments can significantly improve their quality of life and provide them with much greater independence. Resources such as the Internet, word processors, or computer-aided design (CAD) packages are great assets in both educational and working environments but without a suitable interface they cannot be exploited easily.

One of the primary tasks when using a computer is to navigate the on-screen cursor using a pointing device. According to Dennerlein and Johnson, the use of a pointing device accounts for 30-80% of all time spent working at a computer [DJ06]. This requires the operator to accurately position the cursor and maintain stability whilst operating a device switch. Many people with disabilities find traditional input devices difficult to use. According to Hwang et al., symptoms such as tremor, spasm, muscle weakness, partial paralysis or poor coordination can make standard pointing devices difficult, if not impossible, to use [HLKC03]. The traditional mouse used with a computer will often move too quickly for people with uncoordinated movements. As a result, precise manipulations, such as icon selection, can be difficult to perform and will often take a long time. The cursor speed can be slowed down but this means that the mouse physically has to move a much greater distance to perform the same displacement on the computer screen. This movement penalty is clearly undesirable for people with motion impairments. In recent years typical desktop computer screen sizes and resolutions have increased significantly, which means the operator has to move the device even further to reach the extremities of the screen. This can be a major difficulty for people who suffer from fatigue or have a limited range of movement. A higher gain can be provided to the operator but this will reduce the effective width of the targets and make them more difficult to select. Motion-impaired computer users often have difficulty maintaining stability when trying to select small targets. Langdon et al. suggest that difficulties in performing point-and-click tasks often lie primarily in clicking rather than in navigating to the target [LHK+02b]. For example, if the operator experiences a spasm during the button click then this can lead to positional disruption of the cursor. A computer is a highly versatile tool where both software and hardware techniques can be developed to help overcome many obstacles that a disabled person may encounter.

Haptic assistance is the process of using force feedback to aid the operator in human-computer interaction (HCI). This may take the form of guiding the user towards a target or assisting them in its selection. A haptic device offers potential assistance in these areas but so far has not been utilised to its full potential. The concept of using haptic feedback to assist motion-impaired computer users has been addressed in previous studies [KHL+02] [LHK+02b] [HH08]. These techniques have often shown to improve clicking accuracy and reduce targeting times under certain conditions. For example, Keates et al. conducted a number of experiments to determine the effects of haptic assistance using the Logitech Wingman force feedback mouse [KHL+02]. Evaluation was often performed using cursor analysis techniques proposed by MacKenzie et al. [MKS01]. Keates et al. showed that gravity wells can significantly improve throughput and reduce the number of missed-clicks by producing a magnet type effect around the targets [KHL+02]. Targeting times were shown to improve by up to 50% for some operators when using haptic damping. A haptic tunnel network can be applied to an interface to improve the throughput for people that have difficulty following straight paths [LHK+02b].

The results from the literature show that incorporating the sense of touch can significantly improve a person's interaction rates, if implemented carefully. However, these studies only show that traditional haptic techniques, with a static configuration, may aid certain individuals in specific areas [KLCR00] [LHK+02b]. This may not always be appropriate for the disability or the individual, since two people may have the same diagnosis but their level of impairment could vary considerably. As yet the needs of the individual have not been incorporated into haptic interaction for motion-impaired users. There are a number of limitations to traditional haptic techniques that have hindered their inclusion with existing graphical user interfaces (GUIs). Haptic assistance can be intrusive on interaction because it cannot always be easily ignored. This has been most noticeable on 2DOF devices when using gravity wells and high-friction targets because a spring or frictional force is imposed on the operator. These imposing forces can cause difficulties for people with decreased muscle strength (myasthenia) or joint pain (arthralgia). Gunn et al. state that there may be valid reasons for a skilled user to want to ignore the advice provided by a computer system [GMD09]. Asque et al. argue that haptic assistance should not be intrusive on user interaction when used in conjunction with realistic GUIs [ADL11]. A major issue that the current research has not fully addressed is the use of force feedback for multiple on-screen targets in real world graphical user interfaces. Many GUIs do not lend themselves to haptic assistance. Toolbars are often arranged in rows or columns, which causes issues when the cursor has to pass through undesirable neighbouring targets, that contain haptic cues, before reaching the destination. These are often referred to as target distracters and have significantly hindered the inclusion of haptic assistance with existing GUIs. A preferable solution would be to adapt user interfaces specifically for haptic assistance and motion-impaired operators. However, manually designing interfaces for specific techniques or individuals is impractical and not scalable. The aim of this thesis is to produce haptic assistance that can be integrated into existing graphical user interfaces without the shortcomings of previous techniques. The contributing factors that have limited the development of haptic assistance are discussed in greater depth in the following chapter.

The work presented in this thesis has been in collaboration with the Norfolk and Norwich Scope Association (NANSA). The NANSA Adult Centre is based at 200 Bowthorpe Road, Norwich, NR2 3TZ and offers services for youth and adult clients with physical/learning disabilities or challenging behaviour. The adult services are fully committed to empowering people to achieve a smooth transition into work, voluntary placements and community life on completion of their skills development programme. NANSA actively supports and encourages members to access mainstream leisure and special activity groups that match their interests and preferences. Upon successful completion of the skills development programme, clients are supported in taking up work/voluntary placement opportunities through NANSA's close links with The Shaw Trust and Meridian East. Suitable computer access could be very beneficial to the client's personal and professional development. In this study seven participants (3 women and 4 men) with varying degrees of motion impairment were recruited. A number of experiments were undertaken with informed consent once ethical approval had been granted. These were designed to vigorously test and evaluate the haptic assistance within real world interfaces. Two key-workers from NANSA have been actively involved with the project to ensure that the system is adopted by the participants that will benefit the most from it. The feedback provided by the key-workers has been incorporated in the development of the haptic assistance and the final interface. It is anticipated that the results produced in this study will be useful in providing assistance that could significantly improve access to existing computer software.

1.2 Motivations and research objectives

Assistive technologies (AT) have been developed to extend the individual's ability to live independently. These technologies are extremely important for people with physical disabilities, especially because sufferers of cerebral palsy often do not have impaired mental ability. Busby identified that people with a number of different physical impairment conditions have the same desire to use a computer as able-bodied users but have difficulties with most current access systems [Bus96]. Suitable access to a computer interface will provide people with more opportunities in their personal and professional development. Prohibitive costs have often prevented motion-impaired computer users from obtaining these assistive technologies. However, in recent years force feedback devices have become more affordable. The sense of touch can be used to improve human-computer interaction (HCI) but so far has not been utilised to its full potential. Haptic assistance is designed to improve user performance by reducing error rates through physical interaction with the interface.

To physically interact with a virtual environment requires a force feedback or haptic device capable of exerting forces on the operator. The appropriate force calculations that are transmitted to the haptic feedback device are based on interactions between the probe and virtual objects within the environment. This process is commonly referred to as haptic rendering. The objective of this thesis is to develop new haptic rendering algorithms that will enhance methods of interaction for motion-impaired computer users. The motivation behind this research is that suitable computer access will allow people with disabilities to have much greater independence by enabling them to perform tasks that they were previously unable to accomplish, or had great difficulty accomplishing. The majority of graphical user interfaces (GUIs) are designed for standard pointing devices and able-bodied operators. Haptic assistance has the potential to improve access to existing software, which could truly benefit motion-impaired operators and provide them with much greater opportunities. The main research objective is to simplify and improve the efficiency of cursor control through haptic interaction with existing GUIs.

There are many haptic devices available on the market with varying degrees-of-freedom (DOF). The degrees-of-freedom describes the range of movement a device has in terms of its translations and rotations. For example, a traditional mouse can move in the x-y plane and so it has two degrees-of-freedom. Devices that have a three-dimensional workspace and allow translations along the x-axis, y-axis and z-axis, are classed as 3DOF devices. According to Langdon et al., increasing the degrees-of-freedom can improve interaction rates if implemented carefully so that the extra freedom does not over complicate the interface or increase the cognitive workload [LKCR00]. The majority of previous studies have been limited to 2DOF haptic devices and techniques. There are valid reasons for choosing the 3DOF Phantom Omni such as the ability to pass over target distracters more easily [GMT14]. This is discussed in greater depth in Section 2.6.3.

The final motivation for developing haptic assistance is that previous studies have identified that advanced age can make cursor movement increasingly inaccurate [RT96] [SSC99] [KT05]. The general population is growing older and it is estimated that by 2020 almost half the adult population in the United Kingdom will be over 50, with the over 80's being the most rapidly growing sector. As computer usage spreads throughout the population, computer interfaces will have to adapt to meet the needs of the user group. Haptic interaction could have a major influence on this market especially as force feedback devices become more popular and affordable.

The research presented in this thesis has highlighted several significant challenges in the field. Therefore, three significant challenges related to the design and integration of haptic assistance with graphical user interfaces are investigated throughout the remainder of this thesis:

1. Producing techniques that do not require force calibration for optimisation.
2. Alleviating the effects of target distracters.
3. Developing non-intrusive techniques that can be easily used or ignored.

1.3 Novel contributions

In the remainder of this thesis, several significant novel contributions relating to the challenges defined in the previous section are presented:

- A haptic cone technique is proposed to improve clicking accuracy and interaction rates. Haptic cones are embedded into a virtual plane that lies behind the GUI of the selected application. The operator can fall into a cone and will be guided to the apex at the centre of the target. Clamping at the cone apex means that it is very difficult to slip off the target. The technique does not impose forces on the operator and therefore does not require force calibration to optimise interaction for the individual. An additional benefit of not imposing a force on the operator is that the approach will be suitable for a wider range of disabilities. The studies in this thesis prove that haptic cones can provide significant improvements over traditional haptic assistance.

- Deformable cones and deformable switches are proposed to further improve clicking accuracy and optimise interaction with existing graphical user interfaces. Both techniques extend the haptic cone approach to provide assistance that can be easily used or ignored. These deformable techniques almost completely eliminate the effect of target distracters, which have previously plagued haptic assisted interfaces. When using deformable cones, the operator is provided with a flat surfaced plane that they can easily scroll over. If the user requires assistance then they can press into the surface of the virtual plane and the deformable cones will appear. When deforming a cone, the operator is guided towards the apex, which ensures that the click is performed at the target centre. When exiting a target, the cones reform and the restoring force helps guide the operator out. The deformable switches use a similar principle apart from the clicking operation is performed through the haptic simulation of a push-button switch rather than using the device switch. A new haptic rendering algorithm is presented to permit the implementation of these techniques.

- A navigation workbox is presented that combines a rate/position hybrid system to provide access to the whole computer screen for people with a limited range of movement.

- A complete system has been implemented that extracts the GUI features for any software running on Windows and fully automates it with haptic assistance. The system has been designed to allow motion-impaired users to operate quickly and simply, without the need for external assistance or calibration.

1.4 Thesis outline

This section describes the content and arrangement of the remaining chapters in this thesis.

Chapter 2 provides an overview of the literature in the field, focusing on significant contributions to assistive technologies. The chapter identifies a number of difficulties that motion-impaired operators experience when trying to interact with a computer and highlights the potential benefits that haptic assistance has to offer. The current limitations of traditional haptic assistance are assessed to highlight some of the most significant challenges that remain in the field. These challenges are further investigated throughout the remainder of this thesis. A number of cursor metrics have been identified to evaluate pointing device performance during point-and-click tasks.

Chapter 3 describes the implementation of the novel haptic assistance proposed in this thesis. The haptic rendering algorithms have been designed specifically to improve interaction rates for motion-impaired computer users. This extends previous work in the field by developing techniques that can be integrated with existing GUIs more effectively. Many of the shortcomings highlighted in Chapter 2 have been alleviated by utilising the 3DOF capabilities of the Phantom Omni to produce assistance that can be easily used or ignored. A series of new cursor measures are proposed to evaluate the performance of haptic assistance.

Chapter 4 describes the experimental setup that has been used to evaluate the effectiveness of the haptic techniques presented in this thesis. A new experimental task is introduced that uses the Windows on-screen keyboard (OSK) to give a more vigorous evaluation of the assistance in a realistic environment.

The results of the study are presented in Chapter 5. Experiment 1 investigates the performance benefits of haptic assistance that does not require force calibration compared to traditional techniques in a static configuration. Haptic cones and V-shaped funnels are compared against gravity wells and high-friction targets. The focus of the analysis is on error rates and targeting times. Experiment 2 explores the benefits of non-intrusive haptic assistance that the operator can choose to use or ignore. The study concentrates on the evaluation of deformable haptic cones and deformable virtual switches. The analysis is focused on interaction rates and the effect that target distracters have on user performance in a densely populated interface. Experiment 3 uses the ISO 9241-9 task to investigate how target size and shape effect the performance of the deformable haptic assistance. The study is designed to show that the techniques are generalisable for real-world GUIs that contain different target sizes and shape. Experiment 4 investigates the benefits of the haptic workbox in terms of improving computer access for people with severe motion impairments.

The final conclusions are made in Chapter 6. The chapter provides an overview of the work presented in previous chapters and identifies areas in which further work could be continued.

1.5 Terminology

1.5.1 Assistive technology (AT)

Assistive technology provides an interface that helps a person to interact with computer software. Within this thesis the assistance that will be provided is in the form of haptic feedback.

1.5.2 Calibrated systems

A calibrated system uses methods to tune important variables that will maximise user performance. Some methods of haptic assistance require force calibration to optimise interaction.

1.5.3 Declutching

Declutching is the process of disengaging the on-screen cursor, moving the pointing device to a new location and then re-engaging the cursor to extend the range of the device.

1.5.4 Degrees-of-freedom (DOF)

The degrees-of-freedom describes the range of movement a haptic device has in terms of its translations and rotations. For example, a traditional mouse can move in the x-y plane and is categorised as a 2DOF device. Devices that have a three-dimensional workspace and provide translations along the x-axis, y-axis and z-axis are classed as 3DOF devices. Devices that provide these translations and are able to render torques about all three axes have 6DOF.

1.5.5 Force feedback

Force feedback is the term used to describe the interaction between the user and the haptic device. Force feedback conveys real-time information on an object's stiffness, weight and inertia.

1.5.6 Gain

The gain refers to the speed and distance that the cursor moves for a given input of the pointing device. It is often referred to as the device sensitivity.

1.5.7 God-object (Proxy)

The God-object algorithm was originally designed to overcome the problem of object push-through by tracking a history of contact with the virtual surface. The position of the God-object (proxy) is chosen to be the point which locally minimises the distance to the HIP along a surface. A restoring spring force is calculated between the HIP and the proxy using Hooke's law. An example of this is shown in Figure 1.1.

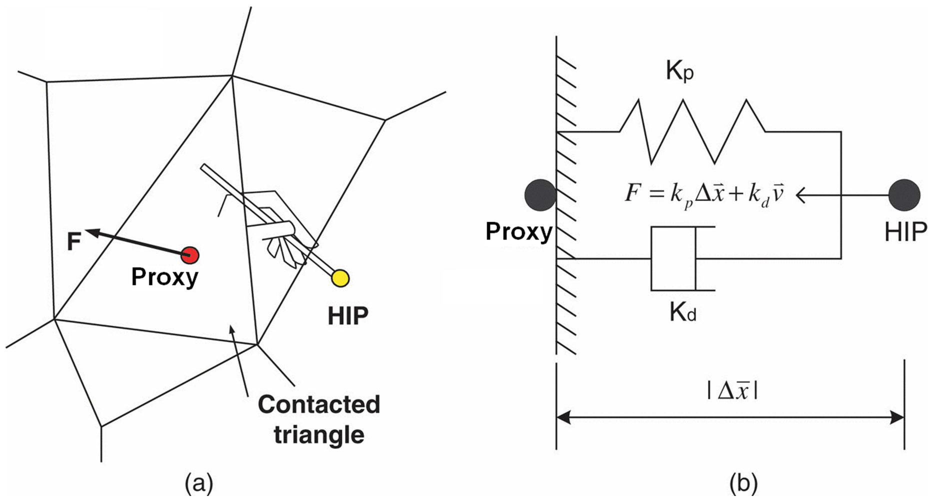

Figure 1.1: The haptic rendering of a surface (adapted from Peng et al. [PZL04]).

1.5.8 Graphical user interface (GUI)

A graphical user interface is a visual type of user interface that allows interaction with computer software.

1.5.9 Haptic feedback

Haptic feedback is the term widely used to include tactile feedback and force feedback.

1.5.10 Haptic interface point (HIP)

When exploring virtual environments, the user will interact with objects through the end point of the probe, known as the haptic interface point. This point represents the virtual representation of the haptic device end effector, as shown in Figure 1.2.

Figure 1.2: The HIP position on the Phantom Omni (adapted from Openhaptics Toolkit Programmer's Guide [OHT14]).

1.5.11 Haptic probe

A haptic probe is the virtual tool that is controlled by the haptic feedback device in the virtual environment. The probe can represent a single point or the virtual representation of the tool in use. It is positioned and oriented in terms of the device translations and rotations.

1.5.12 Haptic rendering

Haptic rendering is the process of computing and generating force feedback in response to user interaction with a virtual environment.

1.5.13 Human-computer interaction

Human-computer interaction is a discipline concerned with the design, evaluation and implementation of interactive computing systems.

1.5.14 Intrusive assistance

Intrusive assistance occurs when a technique is imposed on the operator that they are forced to use and cannot ignore. This can disrupt interaction and have a negative bearing on the overall performance of the system.

1.5.15 OpenGL coordinate system

OpenGL uses the right handed 3D Cartesian coordinate system where the z-axis extends into positive space from the centre of the screen towards the viewer, as shown in Figure 1.3.

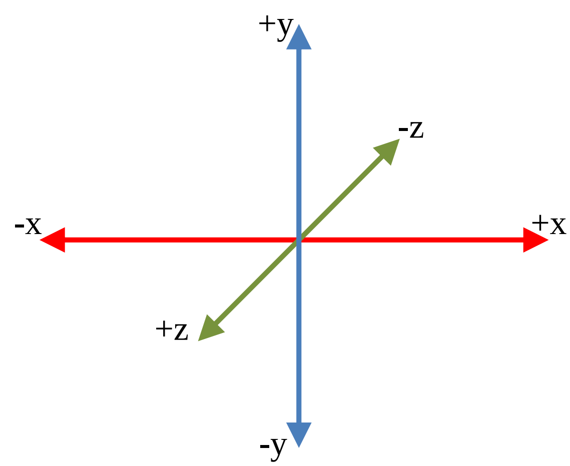

Figure 1.3: The OpenGL coordinate system.

1.5.16 Object push-through

Object push-through occurs when the HIP undesirably passes through a thin mesh. An example of this occurring is given in Figure 1.4.

1.5.17 Surface contact point (SCP)

The contact point located on the surface of a virtual object is often referred to as the surface contact point.

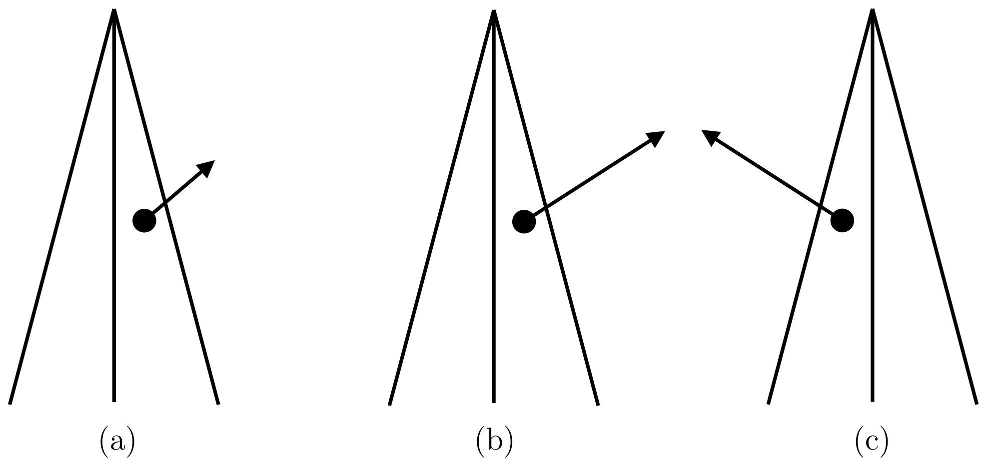

Figure 1.4: Push-through of thin objects. The user touches a surface and feels a small force (a), as he pushes harder he penetrates deeper into the object (b), until he passes more than halfway through the object where the force vector changes direction and shoots him out the other side (c) (adapted from Zilles and Salisbury [ZS95]).

1.5.18 Tactile feedback

Tactile feedback conveys real-time information on the contact surface of an object. For example, heat, pressure, vibration, etc. Interaction with a surface texture is an example of tactile feedback.

1.5.19 Target acquisition

Target acquisition is the process of using a pointing device to position the cursor within a desired target and selecting it by click-and-releasing the device switch.

1.5.20 Target homing

Target homing is the process of navigating the cursor towards the desired target using a pointing device.

1.5.21 Throughput (TP)

Throughput, in bits per second (bits/s), is used to give a measure of the trade-off between speed and accuracy.

1.5.22 Widget

A widget can be considered as a visual element of a graphical user interface (GUI) that performs a function controlled by the user.

Chapter 2

Literature review

2.1 Introduction

In this chapter an overview of the literature in the field is presented, with a focus on significant contributions to assistive technologies. Assistive technology has developed significantly in recent years and has improved computer access for many disabled people that were previously unable to use a computer. There have been a number of areas investigated that have attempted to assist both able-bodied and motion-impaired computer users with human-computer interaction (HCI). Some techniques have utilised haptic assistance whilst others have tuned traditional input devices.

This chapter identifies a number of difficulties that motion-impaired operators often encounter when using a pointing device with graphical user interfaces (GUIs). A detailed background of the current progress of haptic assistance is presented as well as the limitations that have hampered its development with existing software. These limitations are assessed to highlight some of the most significant challenges that remain in the field. These challenges are investigated throughout the remainder of this thesis. Finally, a number of cursor analysis techniques are presented that have been used in previous studies to evaluate pointing device performance during point-and-click tasks.

2.2 Difficulties encountered with cerebral palsy

The following section describes common difficulties encountered with cerebral palsy and associated disabilities. Cerebral palsy is the term given for a series of incurable brain disorders that affect a person's basic motor skills. Approximately one in every four hundred babies born in the United Kingdom has cerebral palsy, which equates to 1,500 people every year [SCO14]. It is important as a software developer to have an understanding of the disability so that areas of assistance can be identified. There are many different types of cerebral palsy and each can affect individuals in different ways. The range of disability for sufferers can vary enormously from mild to severe. This will have a significant effect on the individual's ability to use a pointing device. The list below covers each area and gives a brief description of the condition [NHS14].

1. Spastic cerebral palsy - some of the muscles in the body are tight, stiff and weak due to permanent contraction. Controlled movement can be difficult due to jerky, unpredictable motion. Spastic cerebral palsy consists of these subcategories describing the affected area.

    (a) - Diplegia - Both arms or both legs are affected.

    (b) - Hemiplegia - Either the right arm and leg or left arm and leg are affected.

    (c) - Quadriplegia - All limbs are affected.

    (d) - Monoplegia - One arm or leg affected.

    (e) - Triplegia - Three limbs are affected.

2. Athetoid (dyskinetic) cerebral palsy - Muscle control is disrupted by spontaneous and uncontrolled, involuntary movements. Posture control is also affected.

3. Ataxic cerebral palsy - symptoms include difficulty with balance, shaky movements and speech difficulty.

4. Mixed cerebral palsy - a combination of two or more of the above.

There is no cure for cerebral palsy but there are various forms of therapy that can help a person with the condition. These may include: physiotherapy, occupational therapy, speech therapy, etc. A common misconception associated with people that have cerebral palsy is that they are less intelligent. Learning difficulties may arise as a result of the disability but it is often not the sole cause.

2.3 Problems encountered with human-computer interaction (HCI)

Many disabled computer users find traditional input devices difficult to use. According to Hwang et al., symptoms such as tremor, spasm, muscle weakness, partial paralysis, or poor coordination can make standard pointing devices difficult, if not impossible, to use [HLKC03]. The following section identifies common difficulties that can lead to sources of error when performing pointing device operations. This is followed by difficulties that can arise due to the design of the GUI.

2.3.1 Pointing device operations

The following subsections provide a breakdown of the difficulties encountered by motion-impaired computer users when operating a pointing device. Langdon et al. suggest that difficulties in performing point-and-click tasks often lie primarily in clicking rather than in navigating to the target [LHK+02b]. For example, if the operator experiences a spasm during the button click then this can lead to positional disruption of the cursor. Findlater et al. state that the major challenges motion-impaired computer users encounter with conventional mouse pointing are mostly due to fine corrections at the final stages of target acquisition [FJS+10]. It is for these reasons that the main concentration in this thesis is on the clicking phase.

Single-clicking

The standard mouse requires very fine motor control to position the cursor accurately on a target and maintain stability whilst clicking. In many applications the user is required to click accurately and release the device button whilst still inside the target for the process to execute. This requires a high level of dexterity and coordination. The contraction of muscles when performing a switch press can lead to positional disruption of the cursor, which can result in the operator miss-clicking.

A possible solution for existing pointing devices is to add an external switch, that is separate to the pointing device, to perform the mouse button operations. Therefore, the switch press would not lead to positional disruption of the cursor. The disadvantages of this method are that it requires additional hardware modifications and it deranges interaction due to the separation of the switch from the device.

Double-clicking

Even greater control is required to perform double-clicking operations. Many users have difficulty in maintaining stability during the clicking phase and this can be worsened for double-click operations. The requirement of two successive clicks in a short period of time will often lead to positional disruption of the cursor. For example, Trewin and Pain performed a study on multiple clicking operations and found a 39.5% error rate for the motion-impaired participants when performing double-clicks compared to an error rate of 28.3% for single-clicks [TP99]. Most operating systems (OS) allow the user to adjust the amount of time between the successive clicks but some people can still find the operation difficult to perform.

Many specifically designed pointing devices use a designated button in hardware to emulate double-click operations. The limitation of this approach is that many switches with different functionality can over complicate the interface. Assistive software packages will often provide a virtual mouse where double-click operations are simulated by selecting a widget that then performs the operation after a single-click.

Slipping

Slipping occurs when the operator is unable to maintain a steady cursor between the click and release. For example, the participant may click on a target accurately but slip off it before the operation is complete. Many studies have highlighted this as a major difficulty for both able-bodied and motion-impaired computer users [Bre98] [TKM06] [BSZY11]. According to Trewin et al., slipping whilst clicking is a major source of errors for mouse users with motion impairments [TKM06]. Ideally the cursor would be stationary when pressing the device switch but some users click whilst they are still moving over the target. Another study by Trewin and Pain showed that 28.1% of mouse clicks performed by motion-impaired users occurred whilst the mouse was still moving [TP99].

An assistive feature named “Steady Clicks” was developed to suppress slipping errors by freezing the cursor during mouse clicks [TKM06]. The results produced by eleven motion-impaired participants showed that targets could be selected using significantly fewer attempts. The task completion time improved significantly for five participants with the highest slip rates. A limitation of the “Steady Clicks” approach is that it does not allow the operator to drag-and-drop items, which could be a significant issue for many user interfaces. One of the major benefits of using haptic assistance over the “Steady Clicks” approach is the ability to provide assistance to the operator on-click. When using the “Steady Clicks” approach, the operator needs to have the motor control to accurately position the cursor on the target and click. If they do not have this level of dexterity then freezing the cursor may not benefit them.

Haptic assistance has the potential to assist on-click, during a click and on-release. It is for these reasons that target acquisition techniques are the main concentration of this work.

Accidental clicks

Accidental clicking errors occur if the user involuntarily operates a device switch. This can lead to additional operations, such as pop-up windows, that the person then has to correct. Trewin et al. state that accidental clicks are a major source of errors for computer users with motion impairments [TKM06].

Drag-and-drop

To drag-and-drop an item also requires a high level of motor control. The prolonged contraction of muscles to maintain the switch press can make accurate positioning of the cursor difficult for some operators. A study by MacKenzie et al. concluded that dragging tasks were slower than pointing tasks and that more errors were committed [MSB91]. A useful feature to assist drag-and-drop operations is a technique known as “drag lock”. When using this technique, the operator is required to simply move over the item, click (i.e. press and release a button), move to the new location and click the same button again to drop the item. This functionality can be achieved by including an additional button on the device or GUI that is designated as a “drag lock”.

Physical switch press operation

Each of the previous tasks discussed in this section require the operator to position the cursor accurately inside a target region and press the device switch to perform an operation. It has been observed that some users have difficulty with the physical action of pressing a device switch. For some people this is due to stiffness in the wrist or fingers making it difficult to perform clicking operations. For others that suffer from spasm or tremor the action of pressing the device switch and the contraction of muscles can result in positional disruption to the cursor.

The feedback provided by micro switches on some pointing devices is not always decisive and it can be difficult to determine if the switch has been successfully pressed or not. The surface area of many pointing device switches is often quite small, which can make it difficult for some people to locate them. Section 3.3 presents a new device handle that has been specifically designed and manufactured for motion-impaired users to operate with the Phantom Omni. A haptic device also offers the potential to simulate push-button switches to replace the functionality of those on the device. This concept is discussed in greater depth in Section 2.5.5. The implementation of the novel virtual switches proposed in this thesis are presented in Section 3.4.

Cursor navigation

The on-screen cursor used with a traditional mouse tends to move too quickly for people with uncoordinated movements. As a result, precise manipulations, such as icon selection, can be difficult to perform and will often take a long time. The gain or sensitivity of the device can be adjusted but this leads to a trade-off between speed and accuracy. For example, the cursor speed can be slowed down with a lower gain but this means that the mouse physically has to move a much greater distance to perform the same displacement on the computer screen. For motion-impaired users this additional movement penalty is obviously undesirable. The trade-off between speed and accuracy is closely related to Fitts' law, which is discussed in greater depth in Section 2.7.1. Microsoft have implemented acceleration curves that adjust the gain based on the user input [EPP14]. This aims to allow effective navigation of high-resolution and high-dpi screens whilst maintaining pointer precision at the pixel level. The technique is often referred to as enhanced pointer precision (EPP).

2.3.2 Difficulties with standard GUI design

The design of a GUI can have a direct bearing on how well a motion-impaired computer user can interact with the software. For example, something as simple as the layout can determine whether a person is able to use the interface or not. Graphical user interfaces are typically designed for the average user and it is rare that they will meet the specific needs of people with disabilities. Therefore, the components of that interface may not suit someone with motion impairments. Legislation in recent years has meant that software developers have to provide suitable access to their interfaces. The Disability Discrimination Act (DDA) is an example of legislation passed in the United Kingdom that requires organisations to take steps to ensure their goods and services, including electronic information, are accessible. Although these are positive measures the techniques enforced are still not designed to meet the specific criteria of people with disabilities. According to Bergman and Johnson, users with motion impairments often find it difficult or impossible to use today's common software applications [BJ95]. The most common difficulties with the interface design are highlighted in the following subsections.

Small targets

The majority of GUIs have many toolbars and buttons that perform different operations. As a result, the buttons are often quite small so as to ensure that the majority of the computer screen is available for the main functionality of the software. For example, the ribbons in Microsoft Word are densely populated with small icons, which leaves the majority of the screen available for editing the text document, as shown in Figure 2.1.

Previous studies highlight that motion-impaired computer users often find small targets difficult to select or selection may take a long time [Cas92] [LHK+02b] [KT05]

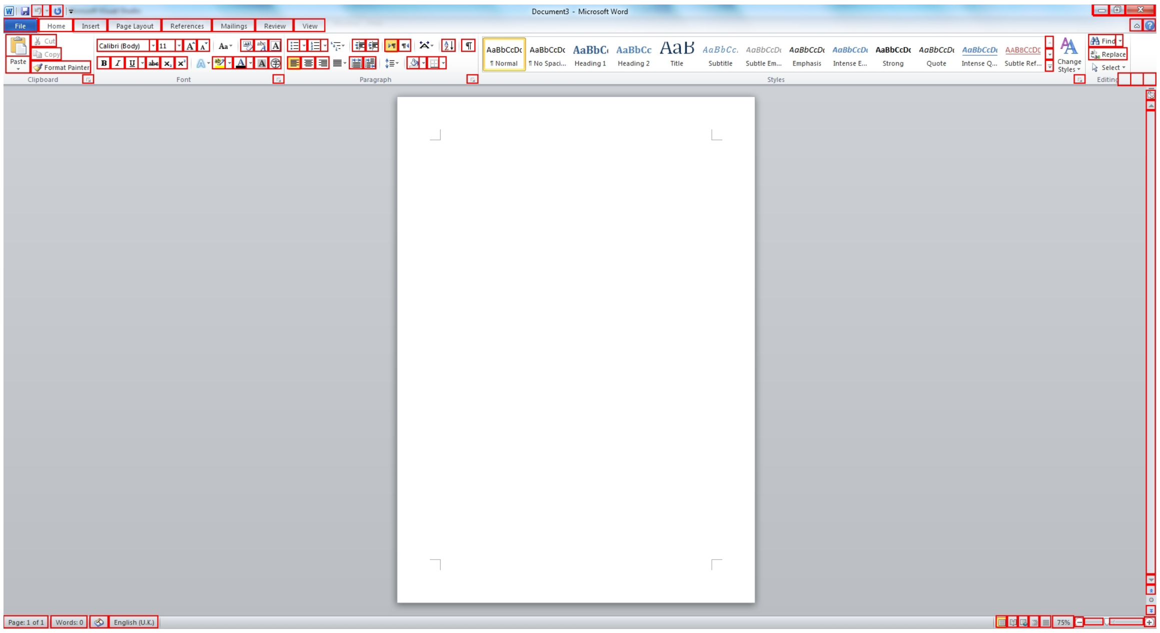

Figure 2.1: The small icons highlighted within a Microsoft Word interface that is displayed on a 15.6” monitor with a resolution of 1920 × 1200.

[CB05] [FJS+10]. One of the most common adaptations is to make targets larger so that they are easier to select. However, commercial software is designed for the average user, which means that the functionality to change the size of the icons for personal preference is often not available.

Studies by Zhai et al. and McGuffin et al. have shown that an effective approach of easing the acquisition of small targets is to only enlarge them when they are required [ZCBLG03] [MB05]. The enlargement is applied to both the visual and motor space. Mandryk and Gutwin state that there are two main problems with techniques that expand a target’s visual space [MG08]. First, the visual changes made are often highly obvious and thus can be distracting, particularly in cases where the expansion is applied incorrectly. Second, visual expansion must either distort the space or occlude nearby areas of the screen [ZCBLG03] [MB05] [HMDH07]. In situations with sparse targets, this may not be a problem; however, if targets are close together or the underlying data is important then the expansion can occlude other objects of interest.

An alternative method is to expand targets in motor space only. A number of studies have investigated a technique commonly referred to as “sticky targets” [WWBH97] [BGBL04] [MG08]. When the cursor passes over a sticky target, the gain of the device is reduced and the cursor moves a lesser distance for a given movement of the pointing device. The majority of studies that have investigated “sticky targets” have been performed using single target aiming tasks, while only a few have investigated the effects of neighbouring distracter targets [WWBH97] [AHL06] [RMI06]. According to Mandryk and Gutwin, one of the limitations of target-aware control to display (CD) gain manipulations is that the system must be able to predict the intended target to be effective [MG08]. For a GUI with a few large icons, target prediction may not be difficult, but for real-world applications with clusters of tightly-spaced, small icons, endpoint prediction is not a trivial task (See Section 2.6.4).

The layout of densely populated GUIs

GUIs are often densely populated with buttons, which can make interaction difficult for people with decreased dexterity. Inaccurate button presses on neighbouring targets can result in undesirable operations that the user will then have to correct.

Gajos et al. suggest that a preferable solution would be to adapt user interfaces to the abilities of the individual [GWW08]. The study by Gajos et al. investigated methods that were designed to improve the performance of motion-impaired users through automatically generated, ability based, interfaces. The system was named SUPPLE++ and models a user’s motor abilities directly from a set of one time performance tests. The use of this system produced results 26.4% faster with the generated interfaces. There were also 73% fewer errors recorded.

Hourcade et al. state that there are limitations of approaches that change the visual design of user interfaces to adapt to a particular user [HNPD10]. The first issue is concerned with learning how to use new software. For example, it would be difficult to write a tutorial for a piece of software if everyone had a different user interface. The second issue arises when working in collaboration with other people. For example, it would be difficult to work in a team if each member had a different GUI layout. Finally, the major shortcoming of this approach is that software developers design the GUI for the commercial market and there is limited flexibility to alter its layout. Therefore, it would be impossible to apply the adaptive technique to existing applications that do not use the same approach to generate the user interface.

Pull-down hierarchies

A common task when using a GUI is steering the cursor through items in a hierarchical pull-down menu, as shown in Figure 2.2. The path that the cursor takes is important. If the path deviates too far from the ideal then a loss of focus can occur and the wrong menu item will be temporarily active. Previous studies from MacKenzie et al. state that such behaviour is undesirable and may impact user performance [MKS01].

Keeping a steady cursor path through combined hierarchical pull-down menus can be a difficult task for people with uncoordinated movements. According to Kobayashi and Igarashi, the typical behaviour of cascading menus tends to cause incorrect selection changes or unnecessary submenu appearance due to straying mouse movement [KI03]. Menu items tend to contain a wide text-based label, which means that the path to a submenu is elongated. A longer or narrower horizontal path reduces the efficiency of mouse operations, especially in tunnel-steering tasks [AZ97]. It is possible to prevent unintentional selection changes of submenus by increasing the delay before they appear or disappear, however, this is another example of the trade-off between speed and accuracy.

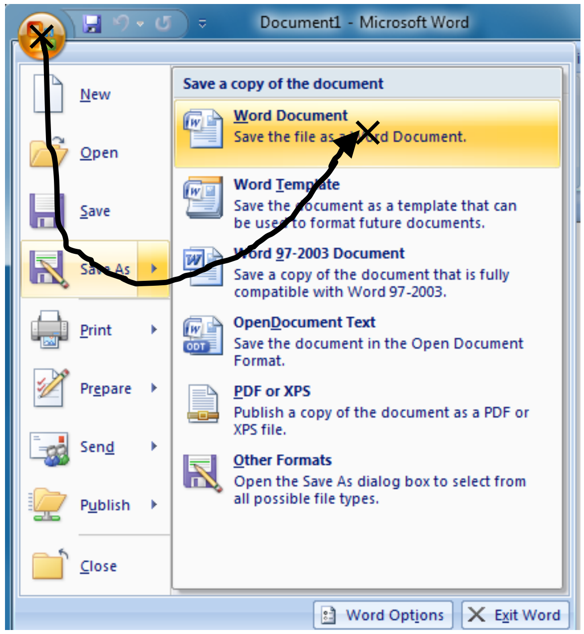

Figure 2.2: Steering the on-screen cursor to save a file in a Microsoft Word interface.

2.4 Alternative input methods for cursor control

People with disabilities are often unable to use a computer because traditional input devices do not meet their needs or commercial solutions are not affordable. The obstacles encountered are different for each individual with a disability. The main difficulty in the communication process for people with motion impairments is the input of data into the system. For example, the use of a mouse or a keyboard demands a high level of accuracy and dexterity. A number of assistive technologies (AT) have been developed to improve computer access for people with disabilities. Some techniques are software driven, whilst others utilise specifically designed hardware. Assistive technologies allow people to have much greater independence by enabling them to perform tasks that they were previously unable to accomplish. However, many techniques often suffer from limitations such as complexity, calibration, affordability, robustness and overall effectiveness. As a result, at least 35% of purchased solutions are never fully adopted [MM99] [RRW00] [KD02] [Daw06]. Dawe reports that the simplicity of the system is of prime importance in the design of input devices and the process of adoption [Daw06]. The following subsections identify some of the most commonly used assistive technologies and their respective benefits and limitations for human-computer interaction.

2.4.1 Speech recognition (SR)

Speech recognition can be a powerful tool for individuals with physical impairments that hinder their ability to use traditional input devices. State-of-the-art speech recognition systems typically provide mechanisms for data entry and cursor control. The most common application is to convert speech-to-text (STT). Two approaches have been adopted for simulating cursor control using speech and these are target-based and direction-based navigation.

Target-based cursor navigation involves assigning speakable identifiers to match the functionality of specific targets on the computer screen. "TalkItMouseIt 2" and "Dragon NaturallySpeaking" are two examples of this type of system [TMI14] [DNS14]. Uttering an identifier will place the cursor over the corresponding target. For example, uttering "save" will place the mouse cursor over the "Save" button. The operator can then click by speaking the function they wish to perform. The main limitation of target-based emulation is that it suffers from layout and usability issues if the number of widgets is high and the interface becomes crowded with many target names.

Direction-based navigation is achieved by moving the cursor in the direction uttered by the user. Karimullah and Sears implemented the continuous movement of the cursor when the corresponding command is uttered, for example, "Move down","-up", "-right" or "-left" [KS02]. The movement comes to a halt upon the "stop" command.

However, participants experienced precision difficulties when stopping because the cursor continued to drift until the speech recogniser had finished processing the utterance. Harada et al. developed a similar technique called "The Vocal Joystick", which assigned each direction a specific vowel sound [HLM+06]. The duration of the sound determines the distance that the cursor travels in that direction.

Lopresti et al. state that while speech recognition is an excellent means of text entry, it is an inefficient replacement for the mouse [LBAG03]. Therefore, its effectiveness may depend on the tasks that the user wishes to perform. Speech recognition systems depend on the user's ability to speak clearly and consistently. This can be a problem for some people with cerebral palsy that have aphasia or individuals that use a ventilator. Even people who speak clearly may require a backup system when working in noisy environments or if their voice is temporarily affected by illness or fatigue. According to Young and Mihailidis, these technical challenges mean that speech based assistive technology often falls short of its potential as an access equaliser for people with disabilities [YM10]. As a result, speech recognition systems are often subject to high abandonment rates [Koe03].

2.4.2 Feature tracking

Feature tracking is a computer vision based system that has been developed to provide computer access for people with severe physical disabilities. The system uses a camera to track movements of a particular body feature (typically the nose) and translates these into cursor movements on the computer screen. An initial calibration phase is performed to evaluate the operator's range of movement. Clicking events are performed based on the cursor "dwell time", i.e. a mouse click is generated if the user keeps the cursor still for a given duration. The "Camera Mouse" developed by

Betke et al. has been shown to improve computer access using consumer level hardware [BGF02]. Twelve people with severe cerebral palsy or traumatic brain injury have used this system and nine of them showed promising improvements, although additional details on training protocol and evaluation methods were not reported. One of the main limitations of the "Camera Mouse" is that the cursor movement is very sensitive to user input at a normal viewing distance, which makes it difficult to perform accurate manipulations. People with insufficient muscle control may find the assistance ineffective. Gain levels can be reduced but this may make the edges of the screen inaccessible. The robustness of current feature tracking systems is limited by challenges associated with recovering lost features and cursor drift. As a result, frequent recalibration is required because of changes in user orientation relative to the camera, involuntary movements, feature occlusion or variations in ambient lighting.

2.4.3 Eye-gaze

Advances in computer vision technology have enabled researchers to develop innovative methods of cursor control. The majority of eye tracking systems consist of a single infrared (IR) camera that tracks the movement of the reflection off the iris for the dominant eye. The cursor is then translated to the location of the user's gaze on the computer screen. An initial calibration phase is often required that consists of looking at multiple reference points displayed successively on the screen. Clicking operations are generally performed using the same dwelling process described for feature tracking. Many systems animate a shrinking frame around the gaze point when the user is fixating on a location. Zöllner et al. investigated blinking versus dwelling and evaluated user satisfaction during a typing task [ZKE08]. The results showed significantly less errors for the blinking method while task completion times did not vary. However, participants rated the fixation method as less tiring and easier to use.

A number of studies have investigated the benefits of eye-gaze technology for people with physical disabilities. For example, Lankford developed an eye-gaze system called "ERICA" that supports text entry and cursor control [Lan00]. A zooming interface was provided to enlarge key features and improve the accuracy of gaze-based clicking. Target selection is accomplished by prolonged fixation on the desired key. A word prediction system was included to speed up typing and reduce errors. Kumar et al. developed a system called "EyePoint" that allows users to perform cursor navigation using a combination of gaze and hotkeys [KPW07]. A progressive two stage refinement process was adopted. During the initial phase the user looks at a certain location on the screen and presses a specific key on the keyboard. The observed portion of the screen is then magnified in a new window. During the second phase the operator looks at the same target within the magnified window and releases the hotkey. A grid of dots is overlaid on top to focus the user's gaze and improve the accuracy of the eye tracking. It was reported that the speed of the gaze-based pointing technique was comparable to the mouse but error rates were significantly higher. Magee et al. developed an eye tracking algorithm called "EyeKeys" using consumer level hardware with video input from an inexpensive USB webcam [MSWB04]. The system uses the camera to track eye movements based on the symmetry between the left and right eye. A further study showed the potential of the system for people with severe paralysis [MBG+08].

Studies have reported that eye-gaze technology is not as precise as the mouse because the size of the fovea restricts the accuracy of the measured point of gaze. Typically, eye tracking systems have an accuracy rating of 0.5 degrees visual angle, or approximately 0.5 to 1cm on a computer monitor at a normal viewing distance. The dynamics of involuntary human eye movements limit the accuracy of eye-gaze technology for fine pointing control, although it can provide effective coarse navigation

[ZMI99]. Input methods that share one channel for control and observation suffer from the "Midas Touch" complex [Jac91]. Such modalities have no intuitive means of differentiating between an input command and other user activity. As a result, errors arise when the system incorrectly interprets user input. A number of studies have reported calibration drift, user fatigue and insufficient range of motion of the eye as factors limiting the effectiveness of eye tracking systems [MR02] [BI03]. Typically, the more robust commercial systems are expensive, which makes them inaccessible to everyday users.

2.4.4 Joystick control

Joystick control is especially useful for people with motion impairments that limit their range of movement. The two common variants are displacement and force joysticks. The output from a displacement joystick is proportional to the stick deflection, whereas the output from a force joystick (isometric) is proportional to the force applied by the operator. The further or harder the user pushes, the faster the cursor moves. This is referred to as first order control, whereas the mouse's mapping of displacement to cursor translation is called zero order control. A joystick allows the operator to navigate the whole of the computer screen with small input movements. Many people with motion impairments are experienced joystick users because they are often fitted to electric wheelchairs to control the speed and direction. The neutral position of the joystick is useful because it acts as a brake, which ensures that the motion comes to rest when the operator releases. Joysticks are especially useful to people with athetoid conditions, such as cerebral palsy, because they tend to be easier to grasp than a standard mouse.

However, previous research has consistently shown joystick control to be slower and less accurate than the mouse [Epp86] [CEB87] [MD96]. A study by Mithal and

Douglas reported that the isometric joystick was 70% slower than the mouse [MD96]. Participants complained that the device was hard to control and they had difficulty getting the cursor to stop accurately inside small targets. The reason for this is that involuntary tremor causes changes in the velocity at which the cursor moves. This makes it difficult for users to precisely position the cursor when attempting to stop at a desired location on the screen. According to Holbert and Huber, a person's disability keeps them at odds with the mouse, leading them to use different devices such as a joystick, which is not as well suited to general interaction with a graphical desktop and application software [HH08].

2.4.5 Trackballs

Trackballs are the preferred pointing device for many computer users with motion impairments [FF01]. For people with low muscle strength, poor coordination, wrist pain, or a limited range of motion, rolling a trackball can be easier than shuttling a mouse across the surface of a desk [WM06]. One of the benefits of a trackball is that they require little space in which to operate. The user can navigate the whole computer screen with small input movements.

Epps compared six pointing devices, including a 4cm trackball, in target acquisition tasks [Epp86]. The results showed that the mouse and trackball were significantly faster than the other devices, but were not significantly different from each other. A follow-up study by Sperling and Tullis reported that the mouse was faster for target selection, dragging and tracing among trackball users [BST88]. Accuracy differences were not significant for target selection and dragging, but were significant for tracing, showing the trackball to be less accurate than the mouse. Further comparisons of pointing devices have shown that trackballs are slower than the mouse when pointing and dragging, and less accurate for dragging [MSB91]. A study by MacKenzie et al. reported that trackballs often move accidentally when clicking inside targets [MKS01]. This is undesirable because a number of studies have identified slipping as a major source of error for people with motion impairments [TKM06] [ADL11]. The effect is more severe when targets are small.

However, trackballs perform relatively well compared to other devices for short straight ballistic movements when crossing a goal. This was the motivation behind the goal crossing system developed by Wobbrock and Myers called "Trackball EdgeWrite" [WM06]. The study showed that using unistroke crossing based gestures for people with motion impairments via a trackball allowed for significantly higher text entry rates in comparison to the on-screen keyboard. The result is a faster and less tedious method of trackball text entry for people who find it difficult to touch-type on a traditional keyboard. One of the major benefits of the "Trackball EdgeWrite" system is that it does not require key-presses, clicking or dragging.

2.5 Haptic assistance

Sections 2.2 and 2.3 highlighted the difficulties that motion-impaired operators experience when interfacing with a computer. The following section justifies the use of haptics as an assistive technology. This is followed by an overview of the most commonly researched haptic assistive techniques and the potential benefits they have to offer motion-impaired computer users.

2.5.1 Why research haptics and the Phantom Omni?

Haptic technology utilises the sense of touch to enable the operator to physically interact with the virtual environment in which they are working. Multimodal interaction is an area of research that has shown that multiple modes of interfacing with a system can improve user performance and satisfaction. The inclusion of haptic feedback in point-and-click tasks has been shown, in several cases, to improve interaction for able-bodied and motion-impaired computer users [OMBG00] [HLKC01] [KHL+02] [LHK+02b] [HH08] [GMD09] [ADL11] [ADL12]. Improvements have been observed both in terms of cursor navigation and target selection. These studies confirm that haptic assistance is an effective method for decreasing the number of missed-clicks and reducing targeting times, the two measures of performance that are most important.

Section 2.3 identified slipping as a major source of errors for people with motion-impairments. Trewin et al. propose the "Steady Clicks" approach to suppress slipping errors but this requires the operator to have the motor control to accurately position the cursor on the target and click [TKM06]. If the operator does not have this level of dexterity then freezing the cursor may not benefit them. Haptic assistance can be used to help clamp a person within a target region so as to suppress slipping errors. Utilising force feedback has the added benefit of being able to aid the operator on-click, during a click and on-release.

There are many haptic devices available with varying degrees-of-freedom (DOF). According to Langdon et al., increasing the degrees-of-freedom can improve interaction rates if implemented carefully so that the extra freedom does not over complicate the interface or increase the cognitive workload [LKCR00]. There are valid reasons for choosing the 3DOF Phantom Omni developed by Geomagic such as the ability to pass over target distracters more easily [GMT14]. This is discussed in greater depth in Section 2.6.3. A 3DOF device also gives the software developer a wider scope to produce assistance that the operator can choose to use or ignore. The importance of this is discussed in greater depth in Section 2.6.2.

This thesis aims to utilise the potential benefits of 3DOF haptic assistance to improve access to graphical user interfaces for motion-impaired computer users. The

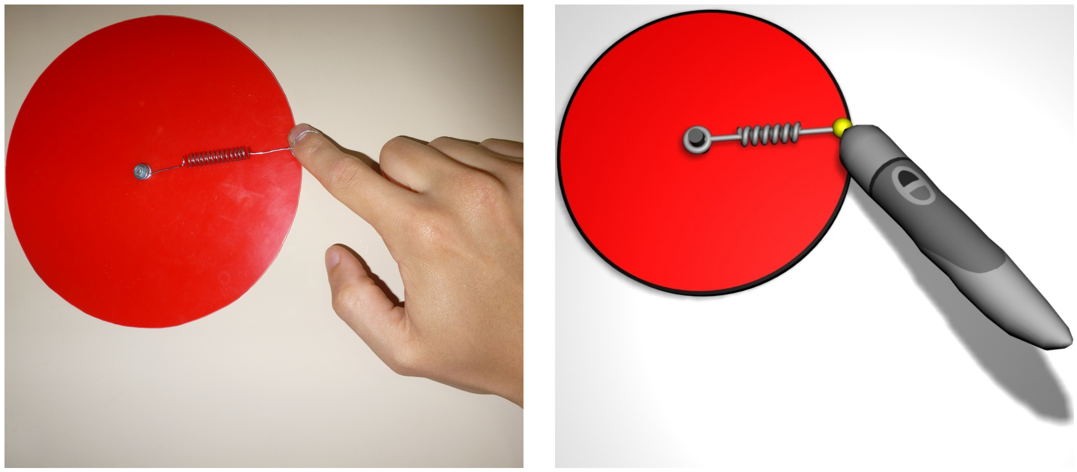

Figure 2.3: When the operator passes over a gravity well, they will experience a spring force that pulls them towards the centre of the target.

following sections identify the most widely researched haptic assistive techniques and the potential benefits they have to offer people with physical disabilities. This is followed by Section 2.6, which identifies the current limitations of haptic research that have hindered its inclusion with existing software.

2.5.2 Gravity wells

A gravity well can be considered as a bounding volume with an inward spring force towards the centre. The concept is presented in Figure 2.3. Hooke's law is used to calculate the spring force to the haptic device. According to Cockburn and Brewster, gravity wells are a useful technique for allowing the operator to more readily select points with the assistance of force feedback [CB05]. They are used to attract the device towards a point location and will typically have some radius of influence. The spring force is designed to clamp the cursor inside the volume until the icon has been selected or the force placed on the device exceeds that limiting the gravity well. The technique extends the pseudo-haptic "snap-to" effect where the cursor is locked to a point of interest [Bie90].

Studies by Keates et al. have shown that gravity wells can significantly improve

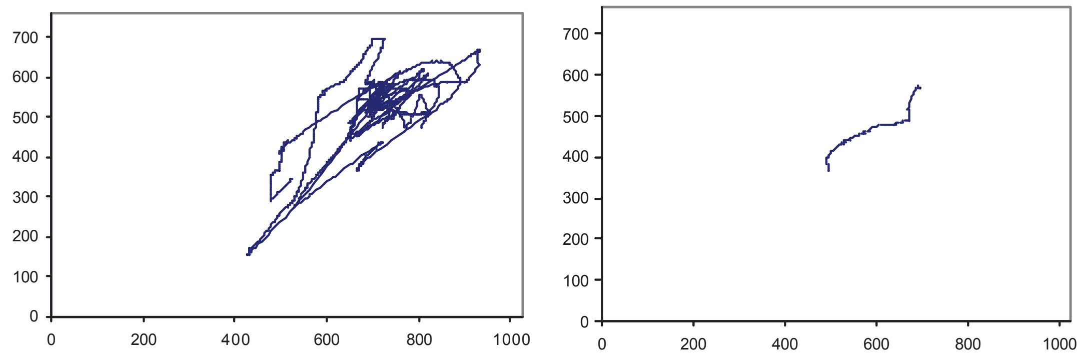

Figure 2.4: A cursor trace for a motion-impaired computer user aiming for a 20 pixel target: unassisted (a) and with gravity wells (b). Image courtesy of [KLCR00].

targeting times and reduce the number of missed-clicks for motion-impaired computer users when performed on an experimental layout [KLCR00] [KHL+02]. The cursor trace in Figure 2.4 illustrates the potential improvements for people with more severe disabilities when using gravity wells in comparison to an unassisted interface [KLCR00].

Gravity wells are the most widely investigated of all haptic assistance. However, there are limitations that have hampered their inclusion with existing graphical user interfaces. The first issue is the effect that target distracters have on interaction [KHL+02] [HLKC03] [HH08]. The second concern is the requirement of force calibration to select a suitable spring stiffness that meets the needs of the individual. These issues are discussed in greater depth in Section 2.6.

2.5.3 High-friction targets

High-friction targets are an extension of the pseudo-haptic technique often referred to as "sticky targets" [WWBH97] [BGBL04] [MG08]. When the cursor passes over a sticky target, the gain of the device is reduced and the cursor moves a lesser distance for a given movement of the pointing device. High-friction targets use a similar

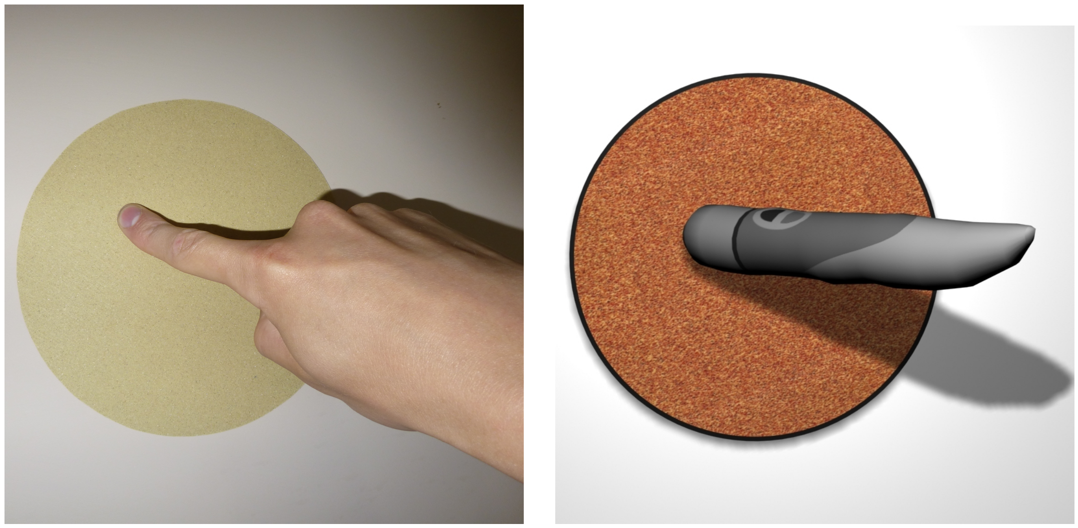

Figure 2.5: When the operator passes over a high-friction target, they will experience a frictional force, similar to passing a finger over sandpaper.

principle except they utilise force feedback to resist the cursor movement. The concept is presented in Figure 2.5. The frictional force is designed to help the cursor adhere to the target so that the operator cannot easily slip off it. Salisbury et al. introduced the concept of using Coulomb friction in haptic interactions because if the user only experiences forces normal to the surface being touched then the sensation of a slippery or frictionless contact is evoked [SBM+95]. Laycock and Day describe smooth and frictionless contact as hindering interaction because the user will often slip off surfaces too easily [LD07]. High-friction targets will be useful in providing stability to people who struggle with the finer movements required for target selection. There are many haptic models that have been created to simulate static and dynamic frictional forces [ZS95] [MRF+96] [RKK97] [LKS02].

However, the high-friction technique presents a practical difficulty because motion-impaired users rarely approach the intended target directly and may inadvertently pass over other targets before reaching the destination. Keates et al. state that if high-friction targets hold the cursor too fiercely then interaction may suffer [KHL+02].

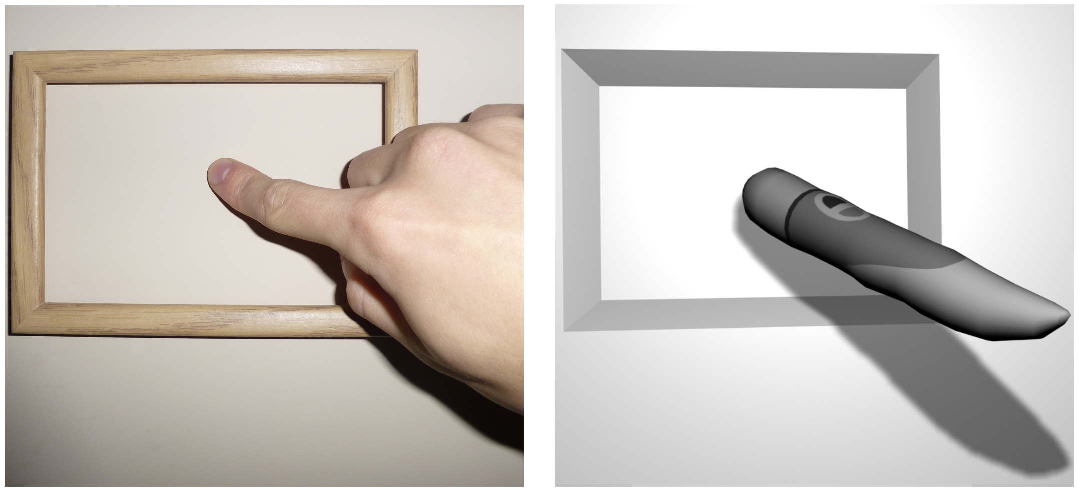

Figure 2.6: When the operator falls into a haptic recess, they will feel the relief in the back of the workspace.

These issues are discussed in greater depth in Section 2.6.

2.5.4 Haptic recess

Oakley et al. created a haptic recess effect where a hole is created in the back of the workspace that has a depth of 2mm and edges sloped at 45° [OMBG00]. An example of the concept is presented in Figure 2.6.

The technique is designed to assist with target selection by allowing the user to fall into the recess. The operator has to make a conscious effort to exit a target by physically climbing the sloped wall. The recess provides stability for clicking and makes it harder to accidentally slip off a target (a problem noted by Brewster [Bre98]). Haptic recesses are one of the few techniques to utilise a 3DOF haptic device to assist with target acquisition. The first experiment conducted by Oakley et al. showed that gravity wells and haptic recesses were the most effective techniques for reducing error rates and decreasing the workload. The results from the second experiment showed that the recess effect provided a significant reduction in the number of times a participant slipped on and off a scroll bar.

The results of this study were only collected from able-bodied users but this non-intrusive technique may benefit interaction for people with motion impairments. The technique does not impose forces on the operator, which means that it can be more easily used or ignored when required. The additional benefit of not imposing a force on the user is that it eliminates the need to calibrate the force level for the individual. The difficulties in calibrating assistive technologies for people with motion impairments are discussed in Section 2.6. The final benefit of the recess approach is that target distracters can be exited easily by lifting off the back of the workspace or by simply climbing the recess wall.

Asque et al. state that the limitations of this approach, in terms of motion-impaired computer users, are that the recess does not assist with guiding the operator towards the centre of a target or with clamping them inside it [ADL11]. Asque et al. go on to argue that providing assistance that clamps the user to the target centre reduces the chances of miss-clicking if the operator slips slightly.

2.5.5 Virtual switches

Section 2.3.1 identified a number of difficulties that motion-impaired computer users encounter when physically pressing a device switch. A haptic device offers the potential to simulate a push-button switch so that the operator does not have to use the one on the device. This concept is presented in Figure 2.7.

The realistic simulation of switches has been addressed in previous studies. For example, Weir et al. developed a "Haptic Profile" concept to measure the subtleties of the physical characteristics of linear push switches [WPC+04]. The system was human actuated and designed to capture a model for the switch properties such as the force, position, velocity and acceleration. The data from this study can be used

Figure 2.7: When pressing a haptic virtual switch, the operator will experience similar feedback to a tactile snap-action switch.

to synthesise realistic haptic sensations. Virtual haptic switches have been implemented by researchers in a number of studies. In terms of hardware Doerrer and Werthschuetzky investigated the force resolution and force-displacement curves using a key-simulator to emulate switches on a control panel [DW02]. Miller and Zeleznik discuss the practicalities of the concept in terms of software design [MZ99]. They describe the force profile of a push-button switch as consisting of an initial springy region where the force increases linearly with displacement, this is followed by a sudden decrease in resistive force and a transition into a "deadband" where the resistive force is constant. Currently there is relatively little literature regarding the use of haptic virtual switches for motion-impaired computer users. This thesis evaluates the effectiveness of specifically designed haptic virtual switches, which are presented in Section 3.4.

2.5.6 Vibrotactile feedback

In recent years vibrotactile feedback has been used extensively in computer game controllers and mobile phones. These devices typically contain a motor that rotates an unbalanced flywheel to produce the vibration effect, as shown in Figure 2.8. A number of studies have investigated the effectiveness of vibrotactile feedback in point-and-click tasks. For example, Akamatsu et al. and Cockburn et al. used tactile vibrations to indicate when the mouse cursor hovered over a target [AMH95] [CB05]. Both studies found that the feedback could improve the performance in certain situations although they reported that the vibration could make users miss small targets.

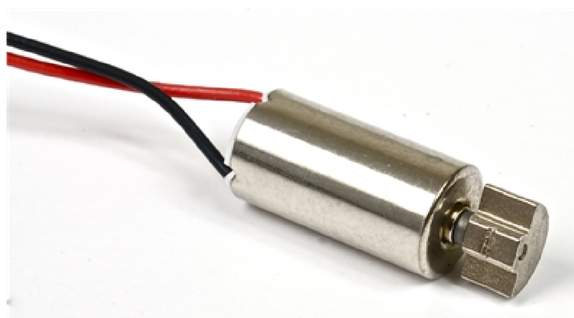

Figure 2.8: A vibration motor containing an unbalanced flywheel. When the flywheel is spun, the operator will experience vibrations from within the pointing device.

A study by Keates et al. concluded that the addition of vibration was a retrograde step for motion-impaired computer users because it almost doubled the time to perform a targeting task [KLCR00]. All of the users expressed displeasure at the vibrating sensation when performing the experiment. These results mirror those obtained using the Phantom with able-bodied users [OMBG00]. Therefore, the research effort will concentrate on alternative methods of haptic assistance.

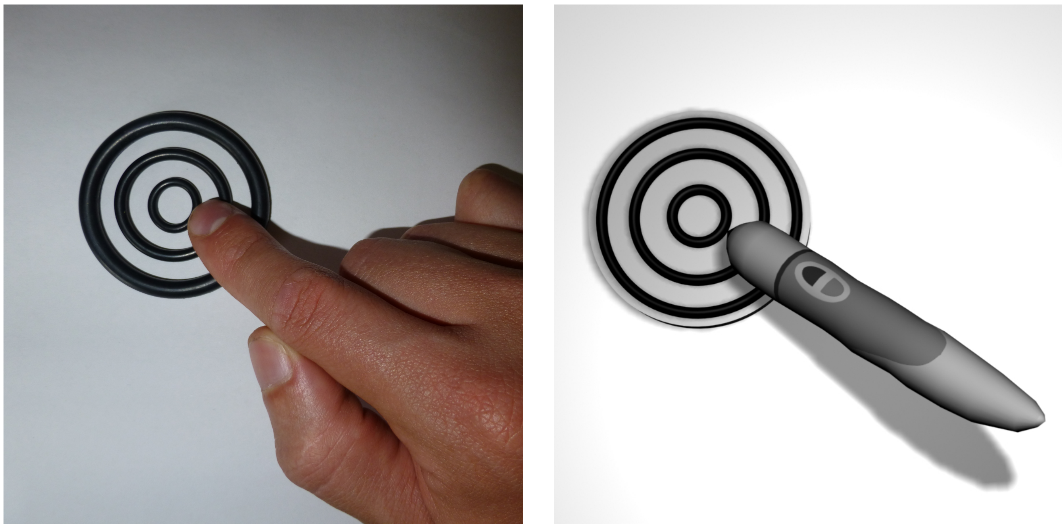

Figure 2.9: When the operator passes over surface textures, they will receive feedback from the ridges, similar to running a finger over nested O-rings.

2.5.7 Surface texture

Force feedback devices can convey texture information by actuating kinesthetic forces on the user's finger, hand or body. Oakley et al. investigated texturing buttons as a potential way of haptically signifying that the cursor is positioned over an object of interest [Oak99] [OMBG00]. The texture used in these experiments was formed of a set of concentric circles centred around the middle of the target, as depicted in Figure 2.9. This texture configuration was chosen because it was felt that it would maximise the possibility that the user would encounter the ridges irrespective of the approach direction. However, the results from the study showed that the texture effect was highly disruptive, with the participants experiencing twice the number of errors that occurred in the control experiment. This demonstrates that not all feedback is helpful and that evaluation should be performed to determine the overall effectiveness of a haptic condition.

Conveying texture information through a kinesthetic force feedback device often relies on much larger forces than those typically experienced on the skin during real texture perception. McGee et al. state that such gross textures can perturb the users' movements so much that the ability to stay on the textured surface is adversely affected [MGB01]. Texture is much more suitable to production by tactile devices such as the "Tractile" developed by Campbell et al. [CZMM99]. Currently, there are no haptic devices on the market that convey both tactile and kinesthetic force feedback.

2.5.8 Haptic damping

People with cerebral palsy often suffer from tremor or spasm, which can make point-and-click tasks difficult to perform. Non-directional viscous damping is a haptic technique that has been investigated to assist in these areas. The technique provides the sensation of moving the device through a viscous fluid and is designed to reduce the magnitude of the spasm or tremor by physically dampening the movement. This ensures that the cursor will not suddenly move across the screen away from the desired target. The concept is presented in Figure 2.10. Millman and Colgate state that damping can help to stabilise the hand as the user tries to achieve a desired position [MC95]. Without this feedback some operators may experience the device as feeling too free or loose and become frustrated when performing small or precise manipulations.

Hwang et al. performed an experiment with motion-impaired participants under four damping conditions: none, acceleration damping, velocity damping and combined damping [HLKC01]. The cursor speed was recorded because it had been noted that a threshold of three pixels/ms could capture the difference between controlled and uncontrolled movement. The results of the study indicated that haptic damping did not significantly improve the targeting times for most operators and that there was little difference between damping types. However, for one participant that

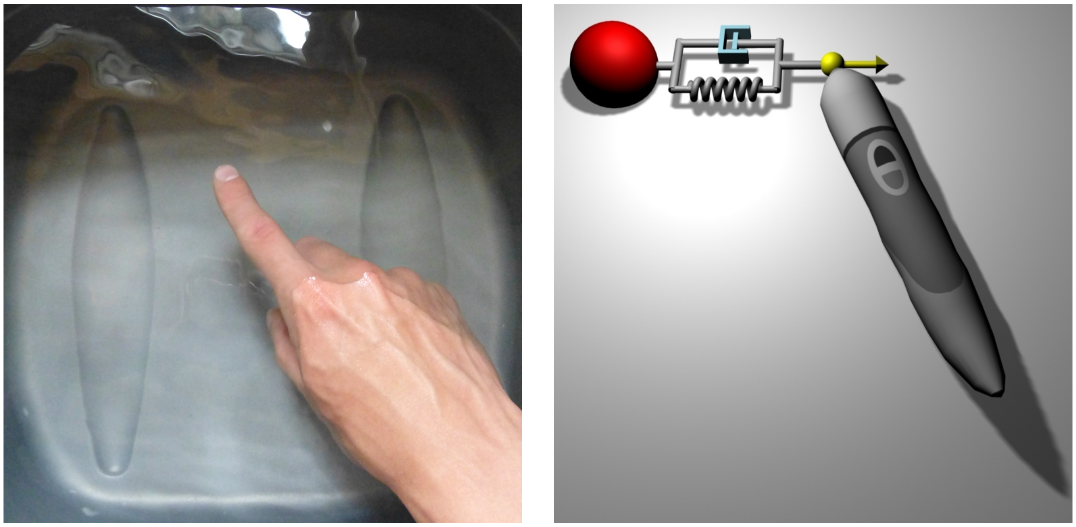

Figure 2.10: When haptic damping is enabled, the operator will experience a resistive force, similar to passing a finger through a viscous fluid.

suffered from spasm the damping technique was shown to reduce the time to reach the target by over 50%. A reduction in the frequency of uncontrolled, high-speed movements of 70-90% was observed for the same participant.

The inherent flaw of this technique is that the operator is under a constant workload due to the damping force also resisting their intended movements. The feedback is intrusive on interaction and the user does not have the option of ignoring the assistance if they wish to. Consequently, an interface can be frustrating to use and will adversely affect user satisfaction. The damping technique may be less useful for other symptoms, such as muscle weakness, because of the extra force required to manipulate the device. An overall performance measure would be required to determine whether damping will be beneficial to a certain user. Determining this trade-off threshold for people with motion impairments is not a trivial task because of the uniqueness of each individual's disability. For example, a certain level of damping may reduce the magnitude of uncontrolled movements but the extra workload could make the device unusable. These issues are discussed in greater depth in Section 2.6.

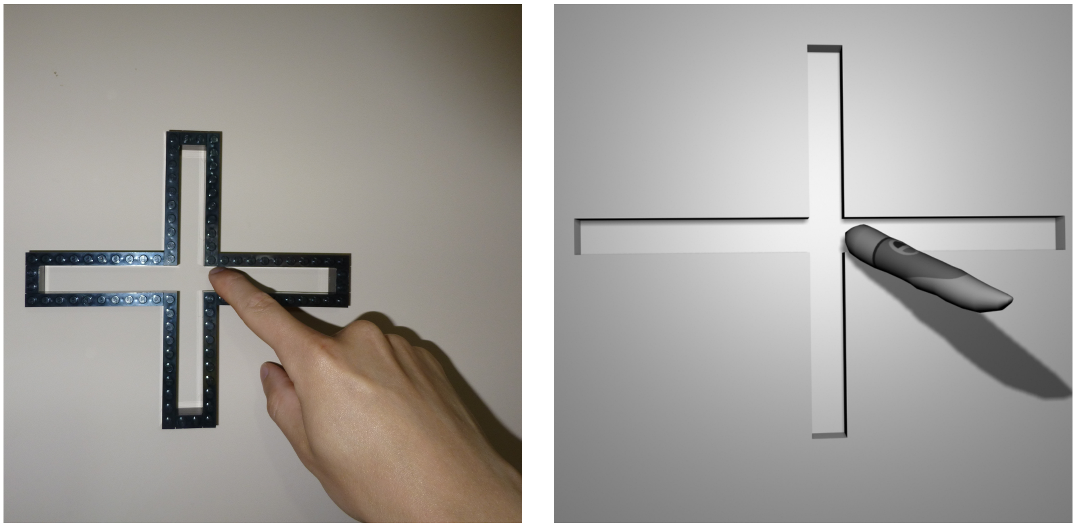

Figure 2.11: The operator can traverse a haptic tunnel network in a similar way to running a finger through a walled maze.

2.5.9 Haptic tunnels

Haptic tunnels have been investigated to improve interaction for motion-impaired computer users that have difficulty following straight paths. They aim to guide the user directly to the target via a tunnel or channel. If the cursor comes in contact with the wall of a tunnel then a restoring force pulls the operator back in to the central channel. The concept is presented in Figure 2.11.

Results from a study by Langdon et al. suggest that haptic force channels may only improve targeting times for sufferers with high degrees of impairment [LHK+02b]. In particular, users with poor navigational ability due to tremor, spasm, weakness and poor control could potentially benefit from forces that prevented movement away from the optimal direction. Langdon et al. suggest that only low levels of improvement were recorded in this study because the difficulties in performing point-and-click tasks often lie primarily in clicking rather than in navigating to the target. For example, spasm during the button click action can lead to positional disruption of the cursor. As haptic tunnels are principally intended to assist the navigation portion of the task, large reductions in targeting times may not be recorded.

The major practicality of the haptic tunnel approach is that the channel should guide the operator to the desired target only. However, to achieve this would require knowledge of which target the user intends to select next. Dennerlein and Yang state that only enabling one force field is an unrealistic simulation for the implementation of force feedback algorithms. If one confidently knew the desired target, why not then select that target automatically without using a pointing device [DY01]? An alternative concept would be to view the tunnels as highways between targets, possibly only to those within a specified radius of the cursor position [KHL+02]. A detailed understanding of the target layout and the preferred routes would be required to implement this technique successfully. For example, if there are a series of buttons that are used regularly but not in the same channel then it could prove a lengthy process navigating the network to get from one target to another.

2.5.10 Literature on the guidelines for designing haptic assistance

A number of studies have highlighted important considerations when integrating haptic assistance into real-world user interfaces. As highlighted in Section 2.6.3 many of these studies have identified target distracters as a major hindrance to haptic integration. Oakley et al. state that the extraneous forces that haptic widgets apply have the potential to alter the paths users wish to take and consequently may reduce their performance and satisfaction [OABG02]. This assertion is upheld in a study that investigated a standard haptically augmented menu system [OBG01]. Oakley et al. propose a number of guidelines that are based on the concept that the force presented should support and not oppose, a user's intent [OABG02]. This entails drawing a balance between allowing users to move where they want as freely as possible and providing forces to improve targeting and reduce errors.

Miller and Zeleznik state that any force feedback applied to a user should be overridable; a user should be able to pop-through, or escape from, any haptically augmented area [MZ98]. Kuber et al. investigated haptic assistance to aid visually impaired Internet users in web page exploration [KYM07]. They state that, ideally, force sensations should be short in duration, perceivable and non-intrusive. Asque et al. propose target acquisition techniques that do not require the operator to oppose a force when exiting a haptic widget [ADL11]. They emphasise that haptic target acquisition techniques should guide the cursor towards the centre of a target so that if the operator slips slightly then they are less likely to miss-click.

It is important to produce techniques that do not impose forces on the operator. The reason for this is that people with decreased muscle strength or joint difficulties may not be able to use the assistance without discomfort. For example, the "snap effect" of gravity wells has been highlighted as a concern for operators with joint difficulties. Similarly, the extra workload imposed by haptic damping has been identified as a concern for people with decreased muscle strength. Techniques that do not impose forces on the user will be suitable for a wider range of impairments.

According to Asque et al., it is desirable to produce haptic assistance that does not require force calibration to optimise interaction [ADL11]. This is due to the fact that a motion-impaired person's needs are not always predictable and so calibration may not necessarily be successful. For example, the operator's needs may change in the short term due to factors such as fatigue or in the long term due to deterioration in impairment. Haptic techniques that do not impose forces on the operator will be suitable for a wider range of impairments. Langdon et al. report that motion-impaired users often exhibit decreased motor control and muscle strength, but not necessarily a decreased haptic sensitivity [LHK+02a]. Studies have determined that the human hand can resolve forces as small as 0.1N [Shi93] [DW02]. Therefore, the force subtleties that the operator experiences need to be carefully considered.

2.6 Current limitations of haptic assistance

A number of studies have been undertaken that have attempted to assist motion-impaired computer users with human-computer interaction (HCI). Given the positive results discussed in Section 2.5 it is surprising that haptic assistance has not been more widely used to improve computer access for people with motion impairments. It is also surprising that haptically enabled interfaces have not been used with commercial software. For example, Claytools allows artists to sculpt 3D models using force-feedback but the interactions that are performed on the GUI, such as 3D Studio Max, can only be performed using the standard mouse [CLA14].

For haptic assistance to truly benefit motion-impaired operators it must be able to be integrated with existing software that they wish to use. The literature has highlighted five major concerns with traditional haptic assistive techniques that have hampered their development and integration into GUIs. These include:

1. Device calibration.
2. Target distracters.
3. The intrusive nature of traditional haptic techniques.
4. The haptic trade-off.
5. 2DOF devices.

These are discussed in greater depth in the following subsections.

2.6.1 Device calibration

LoPresti et al. discuss how each individual's disability is unique and tuning devices to a person's strengths and weaknesses can be critical for success [LBA00]. Several studies have attempted to assist motion-impaired operators with pointing device operations through calibration. Software has been produced that attempts to evaluate an individual's performance and then uses this data to automatically calibrate a number of measures to assist the individual.

Koester et al. investigated the gain settings of pointing devices for users with physical impairments [KLS05]. The gain or sensitivity determines how far the cursor moves on the screen for a given movement of the pointing device. Attempts to configure the gain did not provide a significant increase in performance when compared to the Windows XP default. It was shown, however, that for different individuals the gain did have a significant effect on throughput, percent of error-free trials, cursor entries and overshoot. The cursor analysis was performed using the Compass software package with the Aim test [KLA+03]. Koester et al. discuss that the level of assistance needs to be dynamic because a person's ability can change significantly in the short term due to factors such as fatigue or in the long term due to a progression in impairment.

In 2005 IBM researchers announced a mouse adapter that had been developed to aid users who suffer from hand tremors [LS05]. The device is able to eliminate excessive cursor movement, thereby allowing more normal use of a mouse. This adapter aims to filter out the shaking movements of the hand by utilising a technique similar to that found in the image stabilising systems of some camera lenses. The results of the adapter have shown a much smoother movement of the cursor and significantly improved the accuracy of mouse operations.

A number of studies have been conducted in haptic calibration for able-bodied users. Bayart et al. propose a progressive four stage approach where the level of haptic guidance is gradually decreased to reduce the haptic dependencies on the task [BPK05]. The results showed that full guidance outperformed no guidance with respect to position and improved timing accuracy. However, it also showed that partial guidance outperformed full guidance. This is attributed to the idea that people learn from their mistakes and that if the system provides full control, no mistakes are made and the participant is unprepared for when actual issues arise due to a growing dependency on the system. Li et al. developed a progressive shared control algorithm that exposes participants to an appropriate amount of haptic guidance based on their performance [LHPO09]. The results showed a significant performance increase compared to fixed gain guidance protocols such as shared control and visual fixtures.

The major difficulty in the calibration of traditional haptic techniques for motion-impaired users is deciding the appropriate force levels for the individual. For example, what determines an appropriate spring stiffness for gravity wells or damping level for haptic damping? Each disability is unique to the individual and so a certain configuration may improve interaction for one person but make the interface unusable for another. If an individual's impairment deteriorates then the system needs to detect this and react accordingly, both of which are not trivial tasks due to the uniqueness of disability. The unpredictable data inputted by people with physical impairments can make it difficult to effectively tune methods of interaction [HH08]. There are three important factors to take into consideration with the force calibration. The first is the effect that high force levels may have on the device. For example, Keates et al. discuss that a stiff gravity well may cause the cursor to overshoot out of the other side of a target, due to the physical momentum of the device [KHL+02]. Inappropriate forces imposed by a haptic technique may cause excessive wear that will permanently damage the device. The second consideration is the effects that a haptic condition may have on the operator. For example, some people with joint difficulty may find that the "snap effect" of gravity wells causes them discomfort and therefore they are unable to use the technique. Finally, an interface may contain many haptic targets that the operator has to pass over to reach the destination. If the force levels are too high then the operator may not be able to pass over these target distracters easily. Keates et al. state that if haptic targets hold the cursor too fiercely, then interaction may suffer [KHL+02].

The uniqueness of each individual's disability and the unpredictability of data provided by people with motion impairments poses a major challenge in terms of calibrating appropriate force levels. As a result, this thesis proposes haptic target acquisition techniques that do not require force calibration to optimise interaction. This will eliminate additional calibration variables and alleviate many of the concerns that were presented in this section.

2.6.2 Intrusive haptic assistance

Intrusive haptic assistance occurs when a technique cannot be easily ignored and adversely affects interaction. This has been observed frequently on a 2DOF device when using gravity wells and high-friction targets where a spring or frictional force is imposed on the operator that they cannot ignore. A 3DOF device allows assistance to be ignored, if required, by enabling the operator to lift the tool off the virtual plane, pass over distracters and then re-apply the haptic interface point (HIP). However, this can disjoint interaction if the user has to regularly disengage and then re-engage the device. Non-intrusive techniques are much less constraining on the operator and can be ignored more easily. Gunn et al. suggest that there may be valid reasons for a skilled user to want to ignore the advice provided by a computer system [GMD09]. They go on to argue that force feedback may limit the ability to ignore the advice and therefore be less effective as an aid. Wall et al. discuss the need to produce a suitable level of haptic assistance without constraining the operator too much [WPS+02].

Previous studies have not taken this into consideration when investigating traditional haptic assistance for people that have motion impairments.

Only a limited number of non-intrusive haptic techniques have been investigated. The one that has been most commonly researched is the haptic recess, which was discussed in Section 2.5 [Ram95] [MZ98] [OMBG00]. Oakley et al. produced a recess effect where a hole is created in the back of the workspace, with a depth of 2mm and edges sloped at 45° [OMBG00]. The technique is designed to help with icon selection by allowing the operator to fall into the recess. This provides stability for clicking and makes it harder to accidentally slip off a target (a problem noted by Brewster et al. [Bre98]). The first experiment conducted by Oakley et al. showed that gravity wells and haptic recesses were the most effective techniques for reducing error rates and decreasing the workload. The results from the second experiment showed that the recess effect provided a significant reduction in the number of times a participant slipped on and off a scroll bar.

The data was only collected from able-bodied users but this concept of non-intrusive haptic assistance may benefit interaction for people with motion impairments. The technique does not impose forces on the operator, which means that it can be more easily used or ignored. Target distracters can be exited easily by lifting off the back of the workspace or by simply climbing the recess wall. Force calibration is not required to optimise interaction for the individual, which eliminates one of the many calibration variables that arise when trying to tune a pointing device for people with motion impairments.

2.6.3 Target distracters

Under normal conditions the majority of motion-impaired computer users do not have difficulty with the navigation phase of a point-and-click task [LHK+02b]. However, when haptic cues are placed around icons to aid target selection then the introduction of distracters can disrupt the path users wish to take. Target distracters have continued to be a major concern since the early use of haptics in point-and-click tasks [MZ99] [OMBG00] [DY01] [KHL+02] [HLKC03] [HKLC03a] [AHL06] [HH08] [ADL11] [ADL12]. A target distraction occurs when the cursor has to pass through an undesired haptic cue before reaching the destination. An example of this is shown in Figure 2.12. The force of the haptic cue dragging in the cursor can severely disrupt the route to the target. The effect is amplified if more than one distracter lies along the axis of approach.

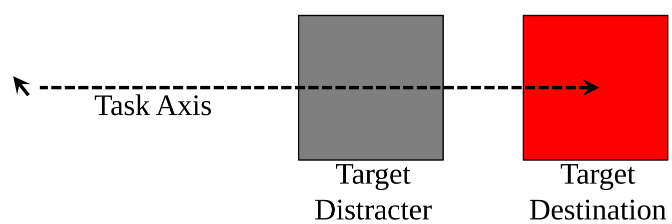

Figure 2.12: A target distracter along the line of the task axis.

Hwang et al. studied the effects of multiple haptic targets on user performance [HKLC03b]. Figures 2.13 and 2.14 show the cursor traces of two motion-impaired users with varying degrees of impairment. It is clear from these figures that distracters surrounding the target can be detrimental to user performance and cause positional disruption to the desired route. The cursor is captured numerous times by the distracter in Figure 2.13, which results in the cursor travelling a much further distance and thus increases the task completion time. Another study by Hwang et al. showed that the layout of distracters surrounding the target can have a significant effect on user performance [HLKC03]. Specifically, the presence of distracters along the task axis in front of the target can be detrimental to performance for some users.

Oakley et al. investigated the addition of force feedback effects to menus where

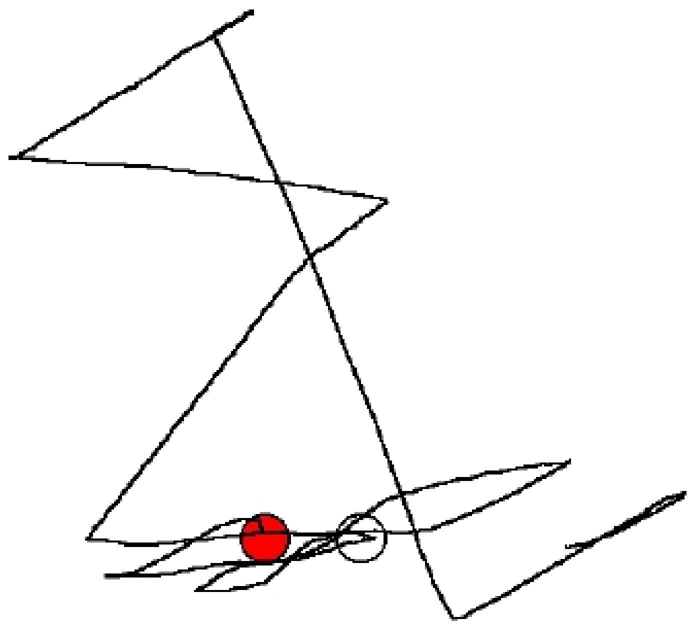

Figure 2.13: The cursor trace of a gravity well distracter along the task axis for a user exhibiting severe motor control difficulties in the dominant hand and arm. Image courtesy of [HKLC03b].

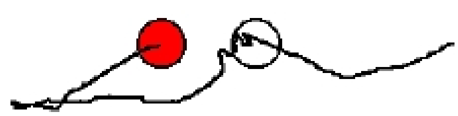

Figure 2.14: The cursor trace of a gravity well distracter along the task axis for a user exhibiting only mild impairment in the dominant hand and arm. Image courtesy of [HKLC03b].

targets are densely stacked vertically [OMBG00]. The results illustrated that a naive application of force could significantly reduce user performance. The fixed forces caused much slower task completion times because users were dragged on to all of the menu items as they moved through the hierarchy.

Previous studies have shown that participants rarely make an effort to avoid target distracters. Consequently, the force that is imposed by a haptic cue will have a direct bearing on the path. Dennerlein and Yang studied several force field configurations in a point-and-click task [DY01]. The results showed that only two of twelve participants actively chose paths to avoid distracters. In most instances the operator continued to move the cursor directly towards the target, despite the high likelihood of getting caught in a distracter and the subsequent effort required to exit it. The majority of participants opted just to plough through neighbouring targets en route to the destination. This may be partly due to the nature of the task and the relatively low penalty associated with getting caught in a distracter. However, for a task with a higher penalty, it is likely that users will tend to make a greater effort to avoid distracters. Hwang et al. state that a motion-impaired user's physical capability may contribute to the apparent lack of adoption of distracter avoidance strategies [HKLC03a]. People with physical disabilities often experience problems with accurate cursor control, which can make it difficult to select a particular target and avoid unwanted areas of the screen. Miller and Zeleznik emphasise that "snap-to" techniques tear holes in the user's input space because it is not possible to specify points near to but not on the snapping point [MZ99].

Some motion-impaired computer users may require higher force levels to provide the necessary clamping that will prevent them from slipping off a target. However, if the interface contains many distracters then this will increase user fatigue due to the operator having to physically oppose the force of the haptic cue before exiting.

If the individual does not have the muscle strength to exit undesired targets then the technique may be unusable. The act of leaving a distracter will often cause the cursor to overshoot, which can impede the next task. Examples of this are observed in Figures 2.13 and 2.14. The effect is often worse with icons in close proximity of each other because the overshoot can land the cursor into an undesired neighbouring target.

Hwang et al. state that target arrangements requiring the cursor to pass through other haptically enabled items can be detrimental to user performance and should be avoided [HLKC03]. For this to be achievable the interface would have to be designed specifically for haptic assistance. Given that most GUIs are designed for standard pointing devices and able-bodied operators it is unlikely that software developers will take this into consideration. As a result, a more effective approach is required to alleviate or reduce the effects of target distracters when integrating haptic assistance with existing interfaces.

In an attempt to reduce the effect of distracters Ahlström et al. created an escape functionality that was designed to help the user exit force fields if the cursor entered an undesired target [AHL06]. The escape functionality deactivates the force after the software has registered six consecutive mouse movements away from the central force point. It is then reactivated once the cursor is moved back toward the target centre. Despite these measures the study concluded that the potential effectiveness of the assistance was still reduced by neighbouring distracters. The results for "sticky targets" showed that the assisted interface provided no significant improvement over the standard one.

A 3DOF device offers the potential to reduce the effects of target distracters by allowing the operator to lift the stylus off the virtual plane, pass over the target distracters and then resume interaction [ADL11]. This approach has been taken for all the haptic techniques discussed within this study. However, if the user has to regularly disengage and then re-engage the device then interaction may become disjointed. It would be preferable if the user was not presented with any target distracters on course to the destination. This could be achieved by giving them the choice of when to use or ignore the assistance.

Cockburn and Brewster state that there has been little research into the use of tactile feedback in more complex and realistic, multiple target displays [CB05]. The issue of target distracters needs to be investigated further so that guidance can be offered to interface designers on how to use effective tactile displays in their systems. Ahlström et al. state that future experiments should focus more on how neighbouring force fields distract the user [AHL06].

2.6.4 Target prediction

Target prediction techniques provide a possible solution to the issues regarding target distracters by only enabling the haptic cues that the operator requires. When navigating towards a target all haptic cues along the approach axis could be disabled so that the cursor is free to move towards the destination. A number of cursor prediction techniques are discussed in the following subsections.

Data based prediction

Data based prediction uses a history of previous inputs to predict the next target. Predictive text and word completion are both examples of data based systems used in modern computing. A similar technique could be applied to a user interface for operations that are often related. For example, if the user selects “Copy” it is likely that this will be followed by “Paste”. The major difficulty with this approach is that a detailed knowledge of the interface or task would be required for it to work effectively. Some tasks may be completely unrelated and so it is unlikely that the software would be able to correctly predict which targets the operator may require next.

Cursor trajectory prediction

A number of studies have investigated target prediction by analysing the cursor trajectory [Mur98] [ASK+05] [LCR07] [HH08]. These techniques monitor the current path of the cursor and use this data to predict the route and distance to the destination. The major shortcoming of this approach is that people do not always follow predictable paths towards a target. According to Mandryk and Gutwin, target prediction may not be difficult for a GUI with a few large icons, but for realistic applications with clusters of tightly-spaced, small icons, endpoint prediction is not a trivial problem [MG08]. Studies that have used trajectory prediction for able-bodied participants have only managed to produce a 75% success rate of correctly predicted targets [Mur98]. It is likely that the rate of correctly predicted targets would be significantly lower for people with motion impairments because they are less likely to take predictable paths. Involuntary movements could significantly distort the accuracy of the predictive algorithm. Holbert and Huber have attempted to use target prediction to reduce the effects of target distracters for haptic cues [HH08]. The targeting performance of the motion-impaired participants improved significantly when the haptic effect was applied to the correct target. Unfortunately, within the study the rate of correctly predicted targets was only 23%, which meant that the overall improvement was unclear. This low performance rate was attributed to the much lower predictability of data produced from people with motion impairments and the sensitivity of the Logitech Wingman in a limited workspace. Dennerlein and Yang state that only enabling one force field is an unrealistic simulation for the implementation of force feedback algorithms [DY01]. If one confidently knew the desired target, why not then select that target automatically without using a pointing device?

Pass-over vs stopping speed

The prediction method discussed in this section compares the pass-over speed vs. stopping speed to only enable haptic cues when the operator requires them for targeting. Oakley et al. investigated the addition of force feedback effects to menus where targets are densely stacked vertically [OMBG00]. The study compared standard visual menus to those with both fixed and dynamically adjusted gravity wells. The dynamic condition adjusted the force level based on the cursor speed and direction. The technique assumes that if an operator is moving rapidly over an item they are unlikely to be targeting it and therefore do not require assistance. The study showed that fixed forces result in much slower completion times due to the participant being dragged on to all of the menu items as they move through the hierarchy. Dynamically adjusted forces significantly reduced task completion times and the subjective workload, as forces were applied only where appropriate. Similar results were reported for tool palettes and desktops [OABG02].

However, it may be difficult to separate these two phases for motion-impaired participants because they often do not perform predictable movements [HH08] [ADL11]. Many people with physical disabilities inherently have slow movement due to the nature of their impairment and so classifying pass-over speed vs. stopping speed may not be achievable. Uncontrolled movements may result in misclassification, which could present the operator with assistance when it is not required and vice versa. Hwang states that velocity based partitioning methods may not always be appropriate in the analysis of trajectories of motion-impaired users [Hwa03]. For example, a spasm may cause the cursor to pass through the target region at a high velocity very early on in the task.

2.6.5 The haptic trade-off

The studies presented in Section 2.5 identify a number of haptic techniques that can improve interaction rates under the right conditions. However, when haptic assistance is introduced to an interface there are often drawbacks associated with the technique that may limit the overall performance. Many of these were identified in Section 2.6. As a result, the level of improvement can sometimes be unclear. This is referred to as the haptic trade-off and is one of the main reasons that haptic assistance has not been successfully integrated with graphical user interfaces. A number of examples are listed below:

- Gravity wells and high-friction targets can assist with selecting small icons but the application of the technique to all buttons introduces target distracters that will often plague the interface.

- A gravity well with a strong spring stiffness may prevent participants from slipping off a target but the physical momentum of the device may cause the cursor to overshoot out of the opposite side of the well.

- Haptic damping can reduce hand tremor but the physical workload will increase due to the operator having to oppose the damping force at all times.

- Haptic tunnels can assist people that have difficulty following straight paths but if the next target is not in the same channel then traversing the network may be a lengthy process.

- Vibrotactile feedback may enhance multimodal interaction but the vibration effect can make small targets difficult to select.

It is clear from these examples that the implementation of haptic assistance requires careful consideration both in terms of interaction rates and user satisfaction.

Some operators may find an interface frustrating to use because the imposing forces may limit the ability to move the cursor freely. Although certain haptic techniques may assist in specific areas, a more general evaluation is required in order to determine its effectiveness on interaction as a whole. The haptic techniques proposed in this thesis aim to alleviate many of the trade-off effects that can limit the overall performance.

2.6.6 Limitations of 2DOF devices

The majority of research that has investigated haptic assistance specifically for people with motion impairments has been performed using a 2DOF device [HLKC01] [KHL+02] [LHK+02b] [HH08]. These devices have a number of limitations that can significantly affect user interaction. The first concern is the inability to ignore target distracters, which were discussed extensively in Section 2.6.3. A 3DOF device offers the potential to reduce the effects of target distracters by allowing the operator to lift the stylus off the virtual plane, pass over distracters and then resume interaction [ADL11].

The second concern is that many 2DOF devices only allow interaction inside of their limited physical workspace [HH08] [ADL12]. For example, the Logitech Wingman has a workspace of just 4cm × 4cm and the user cannot reach objects located outside this area. The consequence of this is that the whole of the computer screen has to be mapped to the confinements of the workspace to allow direct positional control of the cursor. As a result, the cursor gain is often quite high for 2DOF haptic devices, which increases their sensitivity. A large cursor gain will reduce the effective width of the targets and make them more difficult to select. The performance issues that are raised by high gain devices are discussed in Section 2.7.1. Devices such as the mouse do not suffer as significantly from moving larger distances because they have a theoretically unlimited workspace. For example, when the operator reaches the limits of the usable workspace they can lift the mouse off the desk and then put it down on a new location. This is often referred to as declutching. The majority of 2DOF haptic devices are fixed to a platform and so this functionality is not available.

2.7 Cursor analysis techniques

Haptic assistance is designed to improve interaction by reducing error rates and improving targeting times. This section identifies a number of metrics that will be useful in evaluating the effectiveness of the haptic techniques presented in this thesis. The process of selecting a target typically consists of two submovements:

- An initial ballistic phase, which approaches the target.
- A homing phase, which is one or more precise movements to acquire the target.

The cursor measures presented in this section provide a detailed insight into both of these phases. One of the most widely used models in human-computer interaction is Fitts' law, which is discussed in the following subsection.

2.7.1 Factors affecting Fitts' law

Fitts' law is a mathematical model of human motor performance which predicts the movement time (**MT**) from one position to another as a function of the distance to a target (A) and its size (W). The equation is given below. The variables a and b are empirically-determined constants where a represents the start/stop time of the device and b is the speed of the device.

$$\begin{aligned} \mathbf{MT} &= (a+b)\mathbf{ID} \qquad (2.7.1) \\ \text{where } \mathbf{ID} &= \log((A/W)+1) \end{aligned}$$

The definition of Fitts' law implies that targets that are larger and closer together will be easier to select, whereas smaller targets that are further away will be more difficult. These factors have a large influence on the design of GUIs. Studies by Koester et al. have shown that Fitts' law is an appropriate model for people with motion impairments, where larger targets that are closer together can be acquired in less time [KLS06]. Therefore, if it is possible to increase the size of icons or reduce the distance between them then this could significantly improve interaction rates.

Fitts' law can also be used to give a measure of the trade-off between the time taken to select a target and accuracy. This is known as throughput (TP) and is measured in bits/s. The significance of this in terms of the design of haptic assistance is that the technique should aid in either or both aspects of interaction without significantly hampering the other. The equation for calculating throughput is given below.

$$\mathbf{TP} = \frac{\mathbf{IDe}}{\mathbf{MT}} \qquad (2.7.2)$$

$$TP = \frac{\log(\frac{Ae}{4.133 \times SDx} + 1)}{MT}$$

where

$SDx =$ is the standard deviation in selection coordinates along axis of approach.

$Ae =$ is the distance or amplitude of movements.

MacKenzie provides more detailed information on the derivation of these equations [MKS01]. The Phantom Omni used in this thesis is a 3DOF device but the experimental tasks only require target acquisition in the two-dimensional plane of the GUI. The following subsections identify a number of factors that can affect Fitts' based tasks in modern computing.

Screen size and resolution

In recent years the typical desktop computer screen has significantly increased in size and resolution. The consequence of this is that a pointing device has to travel a much greater distance to navigate the whole of a computer screen. This can make a computer inaccessible for people that suffer from fatigue or have a limited range of movement.

The definition of Fitts' law indicates that mouse efficiency has decreased with the increase of screen resolution. For example, Microsoft Word 1.0 was designed for a screen resolution of 640×480 with toolbar icon dimensions of 20×20. However, the same icon in Microsoft Word 2010 may be displayed on a screen with resolutions in excess of 2560×1080. The button size has remained unchanged but it is likely that the cursor will be much further away from the next target than it would have been on a 640×480 display. This can cause difficulties for motion-impaired computer users because they physically have to move the pointing device a much greater distance to reach the destination. This will lead to an increase in movement time (**MT**) and user fatigue. Microsoft aim to overcome these difficulties by adjusting the velocity and acceleration of the mouse cursor based on the user input. The enhanced pointer precision (EPP) algorithm uses acceleration curves to allow effective navigation of high-resolution and high-dpi screens whilst maintaining pointer precision at the pixel level [EPP14]. Koester et al. have shown that EPP is beneficial for many computer users that have physical impairments [KLS06].

Higher resolutions will also result in the icons appearing much smaller on-screen. Screen magnifiers are a useful tool for increasing the visual size of targets but they do not enlarge them in terms of device displacement. Previous studies have investigated techniques such as "sticky targets" to expand icons in the motor space [WWBH97] [BGBL04] [MG08]. This is achieved by reducing the gain of the device when the operator passes over an icon, thus increasing the effective width of the target. The workspace of a pointing device and the cursor gain are closely related to the display resolution on a computer screen. These factors are discussed in the following subsection.

Device workspace and cursor gain

The gain or sensitivity determines how far the cursor moves on the computer screen for a given input of the pointing device. Koester et al. investigated the gain settings of pointing devices for computer users with physical impairments [KLS05]. The results of the calibration did not provide a significant improvement in performance when compared to the Windows XP default. Fitts' law provides the most logical reason for this. An increase in gain will reduce the target distance (A) by reducing the amount of movement required by the device to translate the cursor a given distance on the screen. However, it will also reduce the effective target width (W) and therefore increase the index of difficulty (**ID**). The simultaneous changes in target distance and width tend to cancel each other out, which results in little performance change.

The limited workspace of haptic devices has been identified as a concern when interacting with high-resolution displays or large virtual environments [DLB+05] [HH08] [SHL09] [ADL12]. For example, the Logitech Wingman shown in Figure 2.15 has a limited workspace of 4cm × 4cm. The consequence of this is that the whole of the computer screen has to be mapped to the confinements of the workspace to allow direct positional control of the cursor. As a result, the cursor gain is often quite high for 2DOF haptic devices, which increases their sensitivity. A large cursor gain will reduce the effective width of the targets and make them more difficult to select.

Figure 2.15: The Logitech Wingman force feedback mouse.

The Phantom Omni has a larger workspace of 16cm × 12cm × 7cm but the whole screen still has to be mapped to these dimensions to provide direct positional control of the cursor. Devices such as the mouse do not suffer as significantly from moving larger distances because they have a theoretically unlimited workspace. For example, when the operator reaches the limits of the usable workspace they can lift the mouse off the desk and then put it down on a new location. This is often referred to as declutching. The majority of 2DOF haptic devices are fixed to a platform and so do not have this functionality. A declutching approach could be applied to a virtual plane with a 3DOF device but this would no longer allow the operator to lift the stylus to pass over target distracters.

Previous studies have investigated hybrid rate/position control systems to enable both accurate interaction and coarse positioning of the cursor in large virtual environments (VE). For example, Dominjon et al. created a bubble technique for interacting with large virtual environments using haptic devices with a limited workspace [DLB+05]. When the cursor is located inside the bubble, its motion is position-controlled with a direct mapping to the user's input. When the cursor lies outside of the bubble, its motion is rate-controlled. The operator is able to direct the cursor by pressing against the semi-transparent sphere in the direction they wish to move. Force feedback is provided for interactions between the probe and the sphere. The resultant force is used to govern the rate of cursor movement. The concept of the bubble in a virtual environment is presented in Figure 2.16.

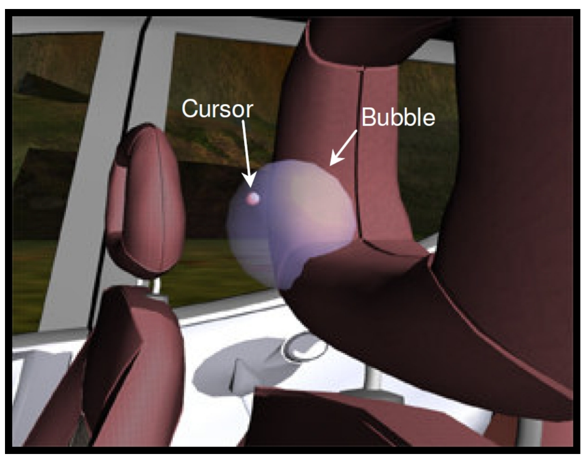

Figure 2.16: The visual display of the semi-transparent bubble in a virtual environment. Image courtesy of [DLB+05].

Casiez et al. created a 2D passive haptic feedback system that consisted of an elastic ring mounted on top of a touchpad [CVPC07]. The technique was designed to allow the user to switch from position to rate control without clutching. Results showed performance benefits when reaching more distant targets, whilst maintaining accurate positional control for precise interactions.

Stocks et al. state that spherical navigation volumes do not correspond well to the workspace of a haptic device and so they propose an automatically scaled navigation cube [SHL09]. The navigation cube is used to translate a protein within the haptic workspace by moving the probe outside of the walls of the cube in the desired direction, as shown in Figure 2.17. The rate of translation is dependent on how far the user penetrates the wall. This approach does not calculate forces between the navigation cube and the probe so as to avoid confusion when interacting with the biomolecule.

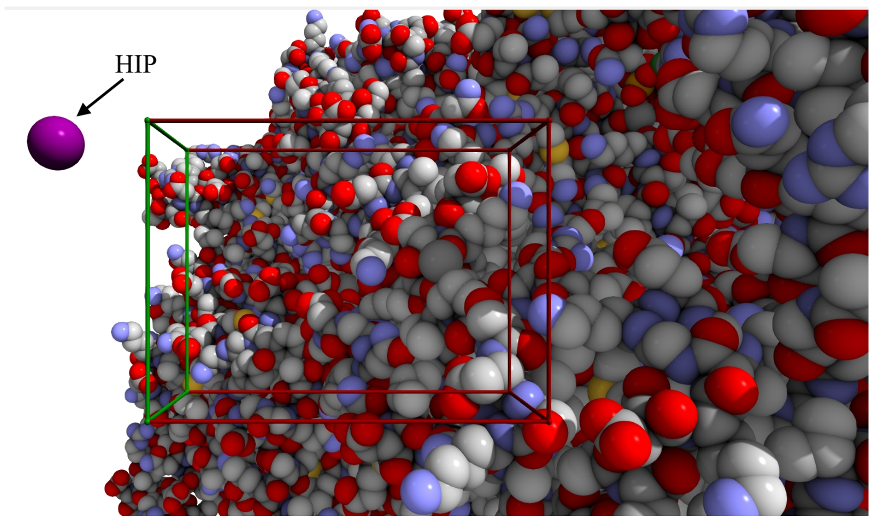

Figure 2.17: As the user moves the device out from the left edge of the navigation cube the protein will start to translate towards the right of the screen.

2.7.2 Movement time (MT)

The previous section discussed how Fitts' law can be used to predict the movement time. The movement time gives a measure for the performance of the participant during a point-and-click task. Haptic assistance is designed to improve movement time but there are occasions where target distracters may be a hindrance and limit the overall performance. This is due to the operator having to perform corrections when passing through and exiting distracters. The movement time will be a useful measure for evaluating the overall performance benefits of haptic assistance and any limitations imposed by target distracters. A comparison between haptic conditions and the unassisted experiment will determine if the newly proposed techniques outperform the traditional ones.

2.7.3 Missed-click

A missed-click is recorded if the click or release (or both) lie outside of the target region. An example of this is shown in Figure 2.18.

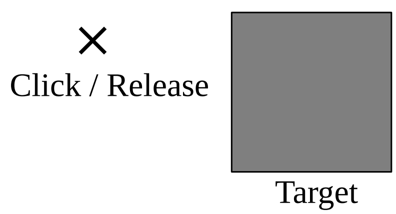

Figure 2.18: A missed-click outside of a target region.

Trewin et al. subcategorise missed-clicks into the following [TKM06]:

1. Near miss - the mouse down position was within 50% of the target radius.
2. Not-so-near miss - the mouse down position was between 50% and 100% of the target radius.
3. Accidental - unintentional clicks, defined as clicks made at a distance > 200% of the target radius, or cases where the user presses down a button and then presses another button before releasing the first one.

2.7.4 MacKenzie's cursor measures

Although Fitts' law indicates that differences exist in movement time and accuracy, it does not give an explanation as to why these exist. If it is possible to understand the reasons why difficulties arise then it will be possible to produce better solutions. The literature indicates that a series of cursor measures proposed by MacKenzie et al. are often used to evaluate pointing device performance [MKS01]. These metrics have also been used in a number of studies to analyse the performance of motion-impaired participants during point-and-click tasks [KHL+02] [MG07] [WG07] [WFL09]. The following subsections introduce each cursor measure.

Target re-entry (TRE)

If the cursor enters the target region, leaves, then re-enters the same region, then a target re-entry is registered. An example is presented in Figure 2.19.

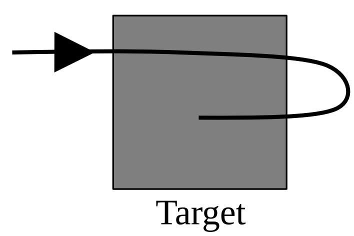

Figure 2.19: The cursor re-entering a target region.

Task axis crossing (TAC)

The task axis is defined as a straight line from the starting position of the cursor to the centre of the target. If the cursor crosses this axis then a task axis crossing is registered, as illustrated in Figure 2.20.

Movement direction change (MDC)

If the cursor's path relative to the task axis changes direction then a movement direction change is registered. This is recorded when the tangent of the cursor path is parallel to the task axis, as shown in Figure 2.21.

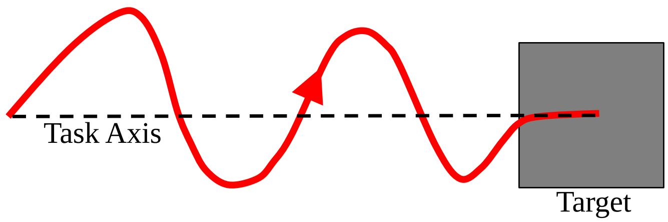

Figure 2.20: The cursor path crossing the task axis when navigating towards a target.

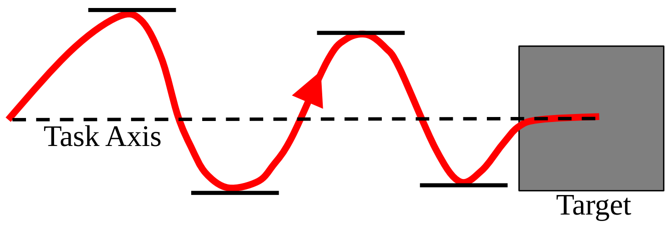

Figure 2.21: Movement direction changes when navigating towards a target.

Orthogonal direction change (ODC)

An orthogonal direction change is registered when the tangent to the cursor path is perpendicular to the task axis, as shown in Figure 2.22.

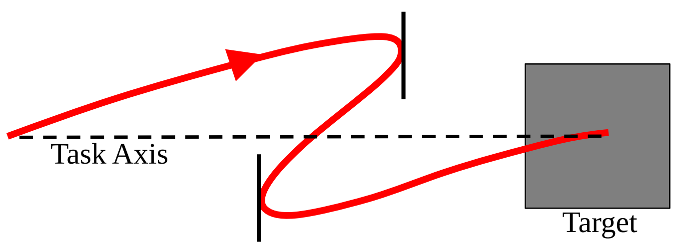

Figure 2.22: Orthogonal direction changes when navigating towards a target.

Movement variability (MV)

Movement variability gives a measure of the extent to which the sampled cursor points lie in a straight line along an axis that is parallel to the task axis. Assuming a task axis of y = 0, then the equation for MV is given below:

$$MV = \sqrt{\frac{\sum(y_i - \overline{y})^2}{n-1}} \tag{2.7.3}$$

where y_i is the distance from a sample point to the task axis.

$\overline{y}$ is the mean distance of the sample points from the task axis.

n is the number of sample points.

Movement error (ME)

Movement error is defined as the mean absolute distance of the sampled cursor points from the task axis, irrespective of whether the points lie above or below the axis line. Assuming a task axis of y = 0, then the equation for ME is given below:

$$ME = \frac{\sum(|y_i|)}{n} \tag{2.7.4}$$

where y_i is the distance from a sample point to the task axis.

n is the number of sample points.

Movement offset (MO)

Movement offset calculates the average deviation of the sampled cursor points along the path from the task axis. It is used to capture the tendency of the pointer to veer left or right during the navigation phase. Assuming a task axis of y = 0, then the equation for MO is given below:

$$MO = \overline{y} \tag{2.7.5}$$

where $\overline{y}$ is the mean distance of the sample points from the task axis.

2.8 Summary

In this chapter an overview of the literature in the field has been presented, with a focus on significant contributions to assistive technologies.

The introduction to the chapter described many of the common difficulties that are encountered by people with cerebral palsy. It is important to have a detailed understanding of the disability in order to produce haptic assistance that will be beneficial to the end user. A number of significant challenges were highlighted in terms of human-computer interaction (HCI). Pointing device operations are often error prone for people with motion impairments and the most common issues were identified throughout the chapter. The layout of the graphical user interface will often have a significant bearing on the individual's ability to interact with the software. As screen size and resolution increases the difficulties for people with motion impairments will continue to rise.

Multimodal interaction is a rapidly growing area of research that utilises more than one of the human senses to interact with a computer. The majority of computer software concentrates on the visual and audio channels. A haptic device allows the operator to use their sense of touch to physically interact with an interface. There are many haptic devices available that are now affordable to a consumer level market. These devices often have a complex hardware implementation that allows for more realistic force sensations. In this thesis the Phantom Omni will be incorporated to provide force feedback to the operator.

A variety of haptic assistive techniques have been explored to improve human-computer interaction (HCI) by reducing error rates and improving targeting times. The majority of motion-impaired computer users have difficulties with target selection rather than cursor navigation and so target acquisition techniques dominate the literature. Gravity wells are the most widely investigated of all haptic target acquisition techniques. A number of studies have shown that they are a useful aid for selecting small targets and reducing missed-clicks. Many haptic techniques have been shown to improve certain aspects of interaction but there are often trade-offs with other areas that may limit the overall performance.

The chapter identified a number of limitations that are currently hampering the development of haptic assistance and its integration with graphical user interfaces. Techniques that require force calibration are not always effective for people with motion impairments because their needs are not always predictable. An individual's habits may change in the short term due to factors such as fatigue or in the long term due to a progression in impairment. Current methods in the literature do not allow the user to ignore haptic assistance and so if the force levels do not suit the individual's needs then this may limit the overall performance benefits. Therefore, it would be preferable to develop haptic assistance that does not require force calibration to optimise interaction. Target distracters have been identified as a major concern in densely populated GUIs because the operator may have to pass over other haptically augmented menus before reaching the destination. The forces imposed by neighbouring haptic cues will often capture the cursor, which can adversely affect interaction rates and user satisfaction. Ideally, the operator would not be presented with any distracters along the task axis when navigating the cursor. Target prediction techniques offer the potential to reduce the effects of distracters by only enabling haptic cues that the user requires. However, predictive methods are not yet accurate enough to use effectively with densely populated GUIs or with people that have motion impairments.

The purpose of this thesis is to produce techniques that are more effective than traditional haptic assistance and not intrusive on user interaction. The aim is to provide the appropriate guidelines to successfully integrate haptic assistance with existing graphical user interfaces. The research presented in this chapter has highlighted several significant challenges in the field. Therefore, three significant challenges related to the design and integration of haptic assistance with graphical user interfaces are identified:

1. Producing techniques that do not require force calibration for optimisation.

2. Alleviating the effects of target distracters.

3. Developing non-intrusive techniques that can be easily used or ignored.

These challenges are investigated throughout the remainder of this thesis. The following chapter describes the implementation of traditional haptic assistance and proposes a number of new techniques to overcome previous shortcomings.

Chapter 3

Implementation

3.1 Introduction

Haptic feedback has the potential to enhance human-computer interaction (HCI) for people with motion impairments. Therefore, a system that allows a user to utilise the sense of touch could prove considerably valuable. The beginning of this chapter describes a number of new cursor measures that have been designed specifically to evaluate the effectiveness of haptic assistance. The remainder of the chapter focuses primarily on the design and development of novel haptic rendering algorithms to allow interaction with existing graphical user interfaces. A number of new haptic assistive techniques are presented that aim to overcome many of the shortcomings highlighted in Chapter 2. The videos presented in this chapter can be viewed directly using the URL or via the accompanying DVD in Appendix A.

3.2 Cursor analysis techniques

The most common evaluation measures are speed and accuracy. Speed is usually reported in its reciprocal form as movement time (MT). Accuracy is usually reported as an error rate based on the percentage of selections outside the target region. These measures can be used to compare the performance of different haptic conditions or participants but they do not give an insight as to why differences exist.

The cursor measures proposed by MacKenzie et al. are often used to evaluate pointing device performance in more detail [MKS01]. Keates et al. also produced a number of cursor measures designed specifically to evaluate motion-impaired computer users [KHL+02]. However, these studies tend to concentrate on the ballistic phase and provide very little insight into the homing phase. This is undesirable given that the majority of difficulties for motion-impaired computer users arise when attempting to select a target [LHK+02b]. The missed-click measure alone is not very useful in analysing what effects haptic assistance have on targeting. For example, if a person miss-clicks without entering a haptic cue then the technique will not have been given an opportunity to assist. The missed-click measure does not provide any indication as to why a missed-click has occurred. Without further analysis there are no guarantees that a haptic condition will certainly help someone that is prone to miss-clicking.

When making a selection in a Microsoft application it is necessary to click and release inside the icon for the process to execute. Many of the previous studies discussed in this thesis do not provide any evidence that this has been taken into consideration. If someone is prone to miss-clicking on-click or on-release then why not adjust the events so that the operation executes on the successful stage? For example, if someone clicks accurately then fire both the click and release events sequentially. The limitation of this approach is that it would be impossible to drag-and-drop items. The cursor analysis techniques proposed in this section are designed specifically to evaluate the effectiveness of haptic assistance during the homing phase for motion-impaired computer users. The new measures will give a greater insight as to why missed-clicks occur and determine the effectiveness of haptic assistance on-click, during a click and on-release. Finally, a number of measures have been proposed that investigate the effects target distracters have on interaction. These will be used to determine whether or not the newly proposed haptic techniques are intrusive.

3.2.1 Missed-click on click

A missed-click on click is recorded if the operator makes a selection outside of the target region but releases the device switch accurately inside. An example is shown in Figure 3.1. The measure is useful for identifying individuals that have difficulty maintaining stability when pressing the device switch but are able to release accurately. A comparison between haptic conditions will identify the techniques that are most effective at providing assistance on-click.

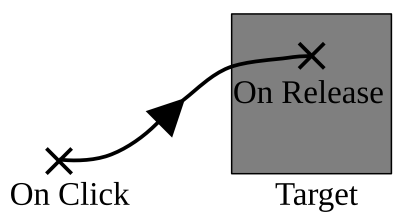

Figure 3.1: A missed-click on click.

3.2.2 Missed-click on release

A missed-click on release is recorded if the operator clicks accurately inside the target but releases outside of it. An example is shown in Figure 3.2. The measure is useful for identifying individuals that have difficulty maintaining stability between the click and release. A major source of error for people with motion impairments is slipping-off a target whilst clicking [TKM06]. The new measure will gauge how effective the haptic assistance is at clamping the cursor within the target during the clicking phase.

3.2.3 Click-release distance travelled

The click-release distance travelled gives a measure of the total distance that the cursor has travelled between the click and release, as shown in Figure 3.3. The

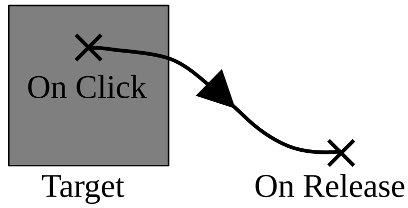

Figure 3.2: A missed-click on release.

measure is useful in determining the stability of the individual during the clicking phase. For example, if the operator moves large distances whilst clicking then it is unlikely that they will be steady enough to select the target accurately. If that individual has a tendency to miss-click on-click (or release) then the click-release distance travelled will give an indication as to why this is the case. The measure will be used to gauge how effective the haptic assistance is at stabilising the cursor during a clicking operation.

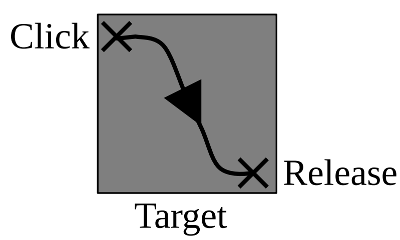

Figure 3.3: The click-release distance travelled.

3.2.4 Click-release displacement

When investigating haptic assistance it is possible for artefacts to be introduced into the cursor analysis. This may be the case for the click-release distance travelled when evaluating stiffer gravity wells. Gravity wells are designed to pull the cursor towards the centre of the target but the device momentum can cause oscillation until the damping takes effect. The oscillation will result in an increase in the click-release distance travelled. The click-release displacement provides an absolute displacement between the click and release and ignores any oscillation artefacts. The measure is presented in Figure 3.4.

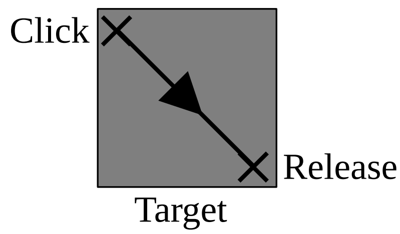

Figure 3.4: The click-release displacement.

3.2.5 On-click distance from target centre line

Slipping whilst clicking has been identified as a major source of error for motion-impaired computer users [TKM06]. If the operator clicks at the centre of a target and slips then it is more likely that the release will occur within the target region. If the click is performed at the edge of a target and the operator slips then it is more likely that the release will miss. The on-click distance from target centre line gives a measure, in millimetres, of how effective a haptic technique is at providing assistance at the target centre. An example is shown in Figure 3.5.

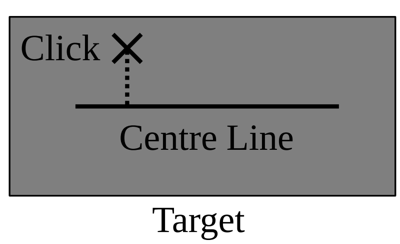

Figure 3.5: The on-click distance from target centre line.

3.2.6 Percentage of experiment time spent on the virtual plane

The amount of time spent in contact with the virtual plane will give an indication of how intrusive each haptic technique is on interaction. If the operator has to regularly lift off the virtual plane to pass over target distracters then this will indicate that the technique is intrusive. The unassisted experiment will give a measure of the natural amount of time spent on the virtual plane. Each haptic condition will be compared against this benchmark. If a similar amount of time is spent in contact with the virtual plane for a given haptic condition then this infers that the distracters are not intrusive because the participant decided to pull off the surface less often.

3.2.7 Experiment distance travelled

It has been observed that the overall distance the pointing device travels will often increase when target distracters are present in an interface. This is due to the operator having to perform corrections when passing through and exiting distracters. The experiment distance travelled will help give an indication of how much disruption is caused by target distracters. The results from each haptic condition will be compared against the unassisted control experiment.

3.3 Device stylus

The Phantom Omni is supplied with a moulded rubber stylus that contains two push-button microswitches. During the preliminary experiments described in Section 4.2.3 it became apparent that some participants were having difficulty physically operating the device switches. For some people this was due to stiffness in the wrist or fingers making it difficult to perform the operation. Other users had issues locating the correct button because their surface area is quite small. The microswitches supplied with the device also lack the desired tactile feedback. They often feel “spongy” and so it can be difficult to decipher whether the switch has made contact or not. These issues detracted from the participants’ concentration during the experiments because their attention was drawn towards operating the device switch rather than performing the task. The stylus is detachable from the device through a 6.35mm stereo audio jack and so it can be replaced to meet an individual’s needs.

A new handle has been designed and manufactured that is better suited to motion-impaired operators. One of the main specifications was to use a lightweight material to ensure that haptic interactions were not affected. Rohacell foam met these requirements because it is a rigid lightweight material that can be easily machined [ROH14]. It is also non-toxic, which was an important health and safety consideration. The new handle was machined on a lathe and a pilot hole was drilled down the centre to mount the stereo socket. The thickness of the stylus was increased to make it easier for motion-impaired operators to grip hold of. The choice of switch was also an important design consideration. The observations from the preliminary experiments indicated that the surface area of the switch needed to be quite large and its tactile response needed to be decisive. As a result, a snap-action tactile switch with a surface area of 10mm $\times$ 10mm was chosen, as shown in Figure 3.6.

Figure 3.6: The snap-action tactile switch used for device switching operations.

Figure 3.7 shows the wiring configuration of the switches supplied with the Phantom Omni. The existing switches are not well located on the device because they are too close together. This often results in the operator accidentally pressing both switches simultaneously. To avoid this issue only a single switch has been connected and mounted to the new handle. The functionality of the second switch will be simulated in software.

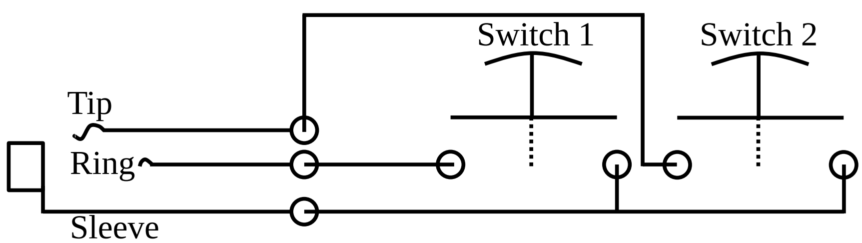

Figure 3.7: The wiring schematic of the switches on the Phantom Omni.

Finally, the foam was sealed using a non-toxic acrylic enamel paint. The finished handle is presented in Figure 3.8.

3.4 Haptic assistance

The software presented in this thesis has been implemented in the C++ programing language utilising the OpenGL API for graphical rendering and the Open Source CHAI3D API for haptic rendering [CBM+12]. The CHAI3D API uses Zilles and Salisbury's God-object haptic rendering algorithm to track a history of contact with a surface [ZS95]. The position of the God-object (proxy) is chosen to be the point which locally minimises the distance to the haptic interface point (HIP) along a surface. The calculated reactive force is proportional to the distance between the HIP and the proxy. The following subsections describe the implementation of the haptic assistance investigated in this thesis. A number of novel techniques are presented

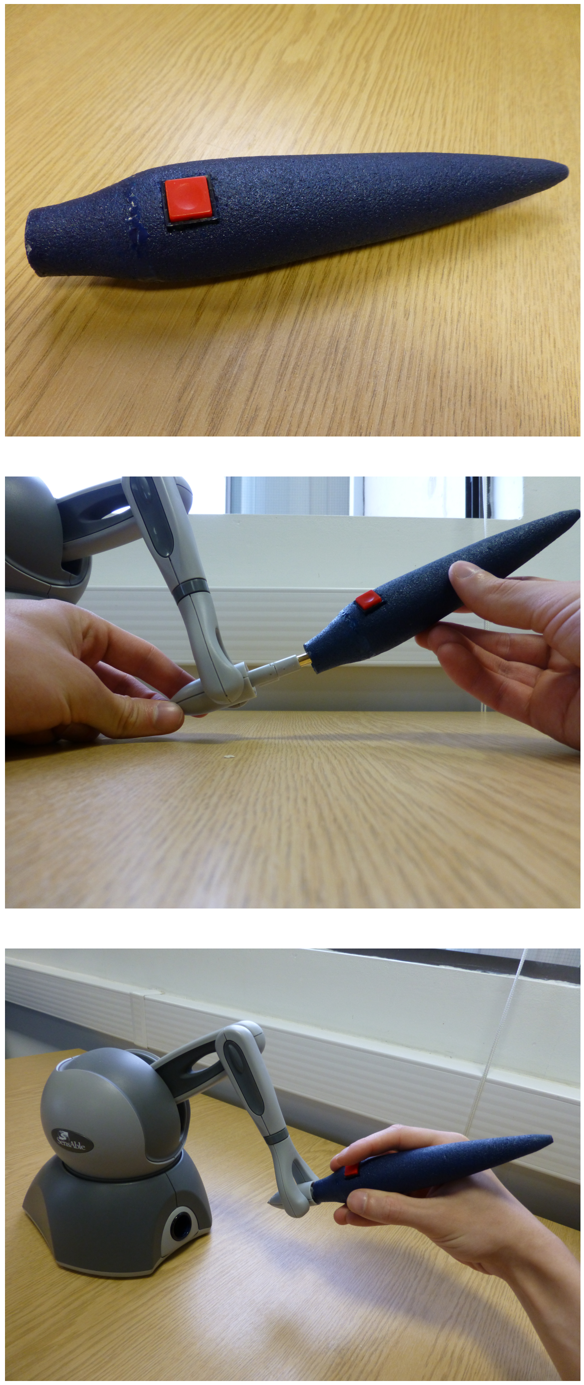

Figure 3.8: Mounting the finished device handle to the Phantom Omni.

that are designed to overcome the limitations of traditional haptic assistance that were identified in Section 2.6.

3.4.1 Gravity wells

A gravity well can be considered as a bounding volume with an inward spring force towards the centre. The spring force is calculated using Hooke's law. When the operator is pulled into a gravity well there are often rapid force direction changes about the target centre that can send the device into oscillation. Consequently, a damping coefficient is often included in the equation to reduce this effect until the motion comes to rest. The equation is given by:

$$\mathbf{f} = k\mathbf{x} - b\mathbf{v} \tag{3.4.1}$$

where $\mathbf{f}$ = resultant spring force

k = spring constant, $\mathbf{x}$ = displacement

b = damping coefficient, $\mathbf{v}$ = velocity of the proxy

Previous studies that have implemented gravity wells have only mapped circular shaped targets. However, this does not correspond well to the square or rectangular shaped icons in most graphical user interfaces (GUIs). The main issue with square or rectangular shaped gravity wells is maintaining a constant force along the edges and at the four corners. For example, if a fixed central pivot was used with a rectangular shaped gravity well then the displacement along each axis could vary considerably. This would result in a much larger force in one axis compared to the other. A minor revision is proposed in this thesis to extend gravity wells to rectangular and square shaped targets.

The approach ensures a constant force along the edges by choosing a moving pivot on the target centre line that locally minimises the distance to the proxy, as shown in Figure 3.9. The resultant force is then calculated based on the displacement between the pivot and the proxy. To avoid force discontinuities at the four corners the displacement is clamped to that of the inner oblique straight oval.

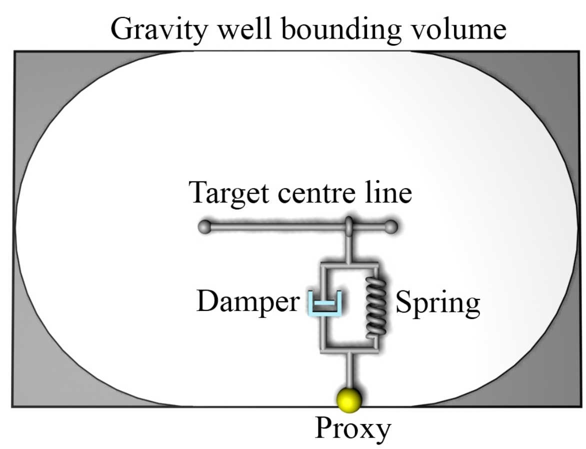

Figure 3.9: The concept of rectangular shaped gravity wells.

Each gravity well is overlaid on top of a virtual plane, which is described in more detail in Section 4.4.2. To allow the cursor to easily pass over target distracters the gravity wells only extrude from the virtual plane by 5mm. This enables the operator to lift the proxy off the plane, pass over distracters and then continue interaction. During the preliminary experiments described in Section 4.2.3 it was observed that the cursor would often overshoot when exiting a gravity well after target selection. To overcome this issue the spring force is disengaged once the target has been acquired. Video 1 demonstrates gravity wells in operation http://youtu.be/exYw0mJo1Hc. Video 2 illustrates the effect of target distracters http://youtu.be/0oKw98qtSnw.

3.4.2 High-friction targets

High-friction targets are designed to aid icon selection by helping the operator "stick" to the target. This is achieved by increasing the friction level of the virtual plane when the operator passes over a target region. As yet very little research has been undertaken to determine the effect of high-friction targets on the performance of motion-impaired computer users. A large criticism of the technique is that the imposing frictional force can seriously hamper interaction if the operator has to pass over many target distracters, especially when using a 2DOF device [KHL+02]. The benefit of using the 3DOF Phantom Omni is that the operator can simply lift the proxy off the virtual plane to pass over target distracters.

A friction model is required to render frictional forces with a haptic device. The method implemented in this thesis is based on Zilles and Salisbury's stick-slip friction model [ZS95]. The algorithm utilises Coulomb friction to calculate a "stiction point" that determines the transition between the sticking and sliding phase. The definition is given below:

$$\overrightarrow{F} = -\mu \|\overrightarrow{F}_n\| \overrightarrow{u}_m \tag{3.4.2}$$

where $\overrightarrow{F}$ is the friction force.

μ is the coefficient of friction.

$\|\overrightarrow{F}_n\|$ is the magnitude of the normal force.

$\overrightarrow{u}_m$ is the unit vector in the direction of motion.

The "stiction point" remains stationary until the tangential force is large enough to overcome the static friction component. When the tangential force exceeds this threshold then the proxy will slip along the surface but will be opposed by a lesser dynamic friction force. The approach is often explained as a spring with a mass attached to it, as depicted in Figure 3.10. The mass will remain stationary until the spring reaches its elastic limit. Once the elastic limit has been met then the motion of the mass is opposed by the dynamic friction force between the two surfaces.

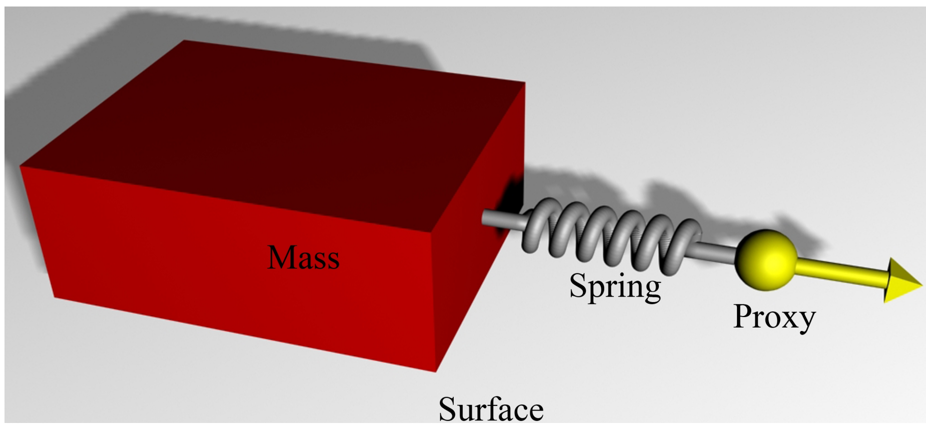

Figure 3.10: The concept of stick-slip friction.

When using high-friction targets, the cursor is projected onto the proxy, rather than the HIP, because it is this object that remains in contact with the surface. Figure 3.11 gives an example of the operator moving along a high-friction surface with the proxy following the HIP. Tracking the proxy ensures that the visual translation of the cursor corresponds to the feeling of resistance through the haptic device.

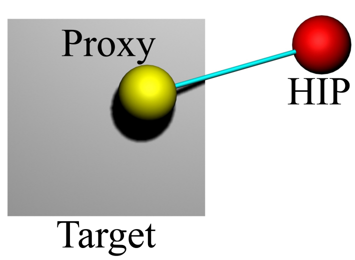

Figure 3.11: The tracking of the proxy and the HIP along the surface of a high-friction target.

Video 3 demonstrates high-friction targets in operation http://youtu.be/-KqBF-KtKl8. Video 4 illustrates the effect of target distracters http://youtu.be/X6hA-tpTU7I.

3.4.3 Haptic cones

In Section 2.6 many concerns were highlighted in regard to traditional haptic assistance. These included the imposing forces of certain techniques and the calibration requirement to optimise interaction. Haptic cones have been proposed in this thesis to overcome these difficulties. A haptic cone is positioned around each target and embedded into the virtual plane, as shown in Figure 3.12.

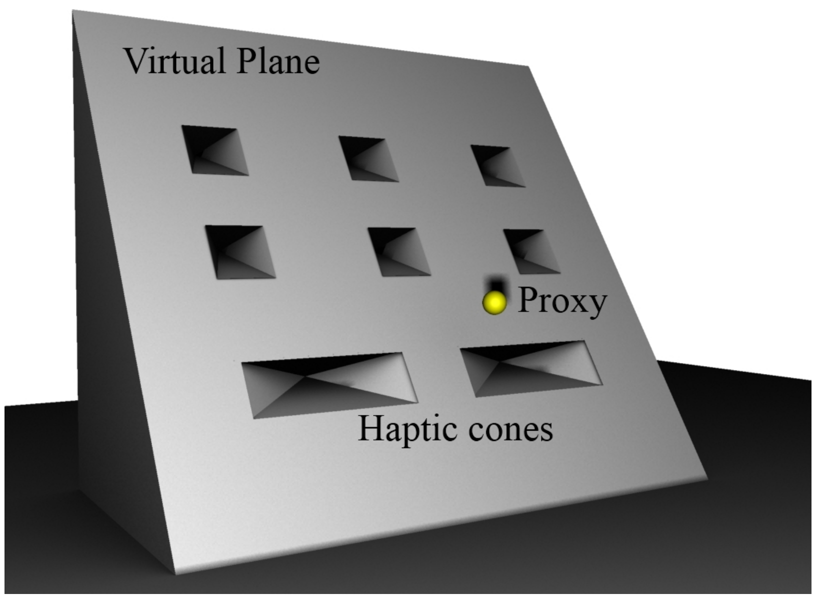

Figure 3.12: Haptic cones embedded into the virtual plane.

Once a cone is entered, the proxy will clamp to the apex at the centre of the target, which provides good stability for clicking. An example of this is shown in Figure 3.13 where the HIP may lie outside of the target region but the proxy remains positioned at the centre. The cursor is projected onto the proxy so that the visual feedback corresponds to that experienced by the user through the haptic device. The number of sides of the cone can be adjusted to correspond with the shape of an interface button. For example, the pyramid shaped cone in Figure 3.13 would be used to map a square shaped icon.

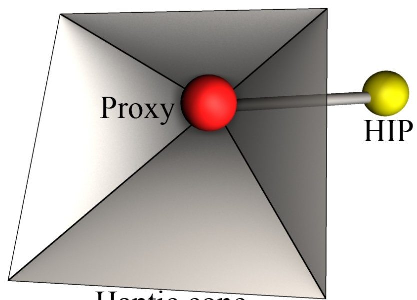

Figure 3.13: The clamping of the proxy at a cone apex with the HIP laying outside of the target region.

Delaunay triangulation has been used to embed the cones correctly within the mesh of the virtual plane [LS80]. The triangulation of the Windows on-screen keyboard (OSK) is presented in Figure 3.14. The depth of the cone is equal to the length of its shortest side so as to provide a suitable slant angle. A wide cone with small depth would have a low slant angle, which would result in the proxy sliding off the apex too easily. Alternatively, a cone with a small width and large depth would require more effort to reach the apex. All haptic interactions are performed using Zilles and Salisbury's God-object haptic rendering algorithm [ZS95].

The haptic cone technique will be suitable for a wider range of users because unlike gravity wells and high-friction targets there are no forces imposed on the operator. Therefore, the approach could be hugely beneficial for people with decreased muscle strength or joint problems. Target distracters can be exited easily by simply climbing the cone wall rather than having to oppose a resistive force. The additional benefit of not imposing a force on the operator is that calibration is not required to optimise interaction. Video 5 demonstrates haptic cones in operation

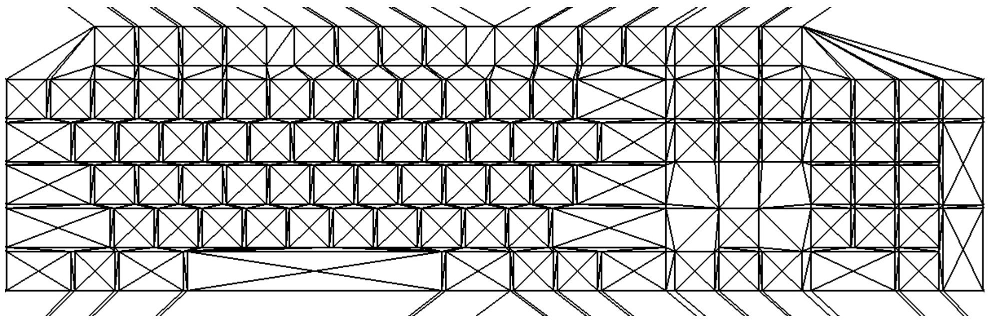

Figure 3.14: Delaunay triangulation of the Windows OSK for pyramid shaped haptic cones.

http://youtu.be/N6N61ZA9Zxk. Video 6 illustrates the effect of target distracters http://youtu.be/cytiXh3aXFs.

3.4.4 Haptic funnels

The V-shaped funnel proposed in this thesis is designed to guide the operator towards the centre of a target. When the cursor lies outside of a target region, the funnel is rotated so that the two walls are always facing towards the proxy, as shown in Figure 3.15(a). When the target region is first entered, the funnel is clamped to its current orientation, as shown in Figure 3.15(b). The operator can then use the funnel walls to guide the cursor to the centre of the target. The joint between the two walls helps clamp the proxy to the centre, which provides additional stability for clicking.

Each funnel is overlaid on top of the virtual plane and extrudes by 5mm. This ensures that the operator can easily pass over undesired targets by lifting the proxy off the virtual plane. The technique does not impose a force on the operator, which means that calibration is not required to optimise interaction. Target distracters can be exited easily by leaving the funnel in the opposite direction to which they were first entered. All haptic interactions are performed using Zilles and Salisbury's God-object haptic rendering algorithm [ZS95]. Video 7 demonstrates haptic funnels in operation

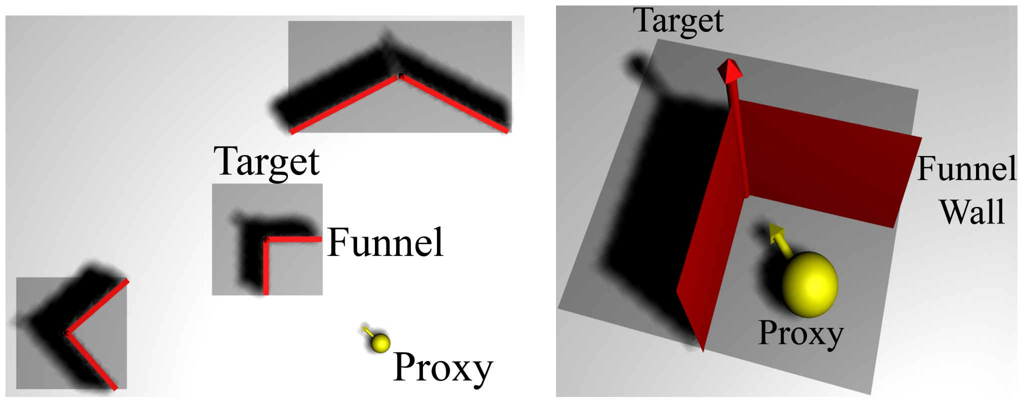

(a) (b)

Figure 3.15: The funnel walls are orientated towards the proxy as the user approaches the target (a). The funnel orientation is clamped once the proxy has entered the target region (b).

http://youtu.be/eQkrul0vhkg. Video 8 illustrates the effect of target distracters http://youtu.be/rvxUPBp_z4M.

3.4.5 Deformable cones

One of the difficulties in the development of haptic assistance has been providing the operator with techniques that they can choose to use or ignore. External switches or gestures that enable and disable haptic cues can disjoint interaction and are not always intuitive. The haptic cone technique proposed in Section 3.4.3 has been shown to improve clicking accuracy and throughput [ADL11]. The approach provides effective clamping at the centre of a target, which reduces the likelihood of slipping off it. Distracters can be exited more easily by simply navigating the cone wall. However, it is still difficult to smoothly scroll over the virtual plane due to the cones embedded in it. Deformable haptic cones have been proposed as an extension to allow the user to choose when they require assistance and when to ignore it.

The virtual plane is embedded with deformable haptic cones that emerge when

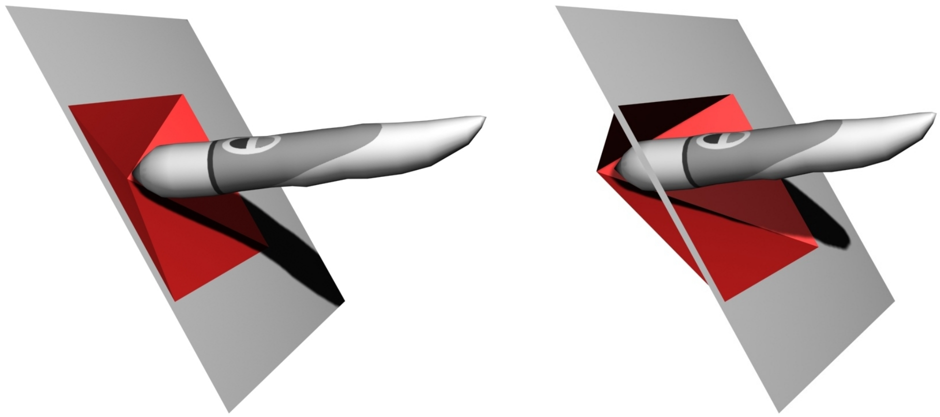

(a) (b)

Figure 3.16: The deformation of a haptic cone by the virtual tool (a). A fully deformed haptic cone (b).

the operator presses into the surface, as shown in Figure 3.16(a). When deforming a cone, the proxy is guided towards the apex, which provides good stability for clicking. The maximum depth limit of the cone is equal to the length of its shortest side so as to provide a suitable slant angle. An example of a fully deformed cone is shown in Figure 3.16(b). A relatively stiff surface is required to support the user's arm whilst scrolling across the screen and to ensure that the cones do not deform too easily.

A new haptic rendering algorithm has been created for deformable cones because the standard God-object approach does not handle both rigid and deformable objects. The implementation is presented in Algorithm 3.1. The cones need to be deformable in the sense that they emerge from the virtual plane but they also need to have rigid sides so as to guide the operator towards the centre of the target. The difficulty with this approach is that two forces require calculation, i.e. the restoring force of the cone apex to the surface of the virtual plane $\widehat{F}_1$ and the restoring force of the HIP to the proxy on the cone surface $\widehat{F}_2$, as depicted in Figure 3.17. It is not possible to simply sum $\widehat{F}_1$ and $\widehat{F}_2$ because the resultant force would exceed that capable of the Phantom

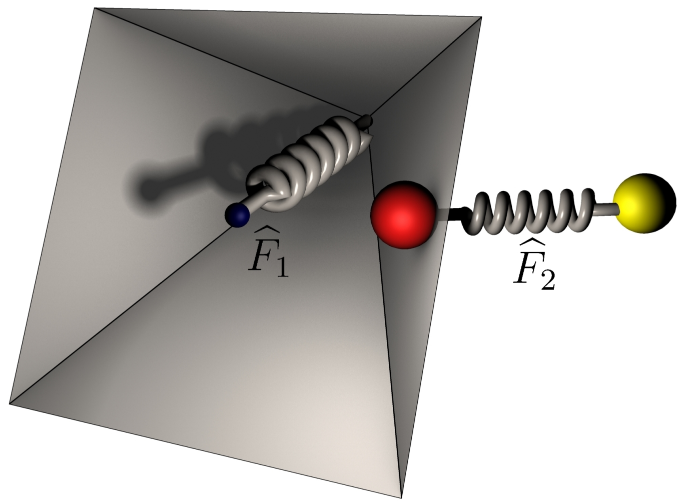

Figure 3.17: The two force calculations required to restore a deformable cone.

Omni. Summing the two forces would also cause a design conflict because the force experienced by the operator needs to be continuous when transitioning between the virtual plane and deformable cones.

To overcome this problem the force rendering algorithm mixes the magnitude and direction of the two computed forces. The initial phase calculates the force direction based on the vector between the HIP and the proxy but the magnitude of the restoring force is governed by the depth of the cone apex in relation to the surface of the virtual plane. This ensures that the magnitude of the restoring force at a given penetration depth of the virtual plane is the same when deforming a cone, as depicted in Figure 3.18. Therefore, the operator will not experience any force discontinuities when transitioning between the two surfaces.

If the x or y components of $\widehat{F}_2$ exceed those calculated previously then they are used as the new resultant force. This ensures that the cone walls have sufficient stiffness to provide effective clamping at the target centre. The traditional God-object implementation can suffer from pop-through when the proxy is in contact

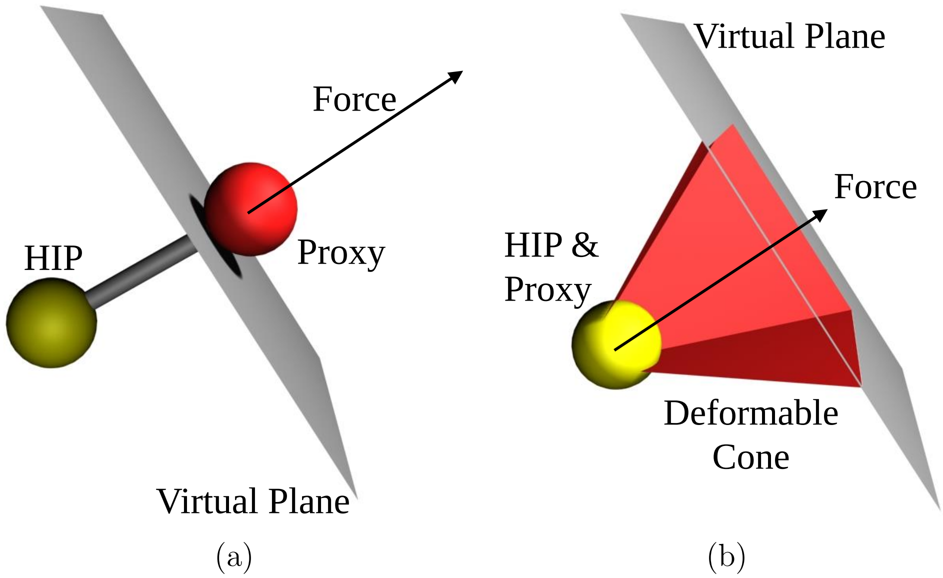

Figure 3.18: The restoring force of the virtual plane (a). The restoring force at the apex of a deformed cone (b).

with a mesh that has moving vertices or geometry transformations. To avoid this issue a constraint has been added that ensures the proxy remains on the correct side of a surface.

Finally, force shading is applied to the edges of the cones to ensure that the proxy does not "catch" when passing over potential distracters. This is necessary because the operator will inevitably deform a cone slightly when scrolling over the interface. The catching effect could impact user satisfaction and disrupt the path of the cursor. No force shading is applied at the cone apex because it is the well defined edges that provide effective clamping at the target centre. The force shading is achieved using spherical linear interpolation (SLERP) between the previously calculated force vector, $\widehat{v_1}$, and the normal of the virtual plane, $\widehat{v_2}$. The threshold, u, of the SLERP is governed by the distance of the proxy from half way up the cone wall to its shared edge with the virtual plane, as shown in Figure 3.19.

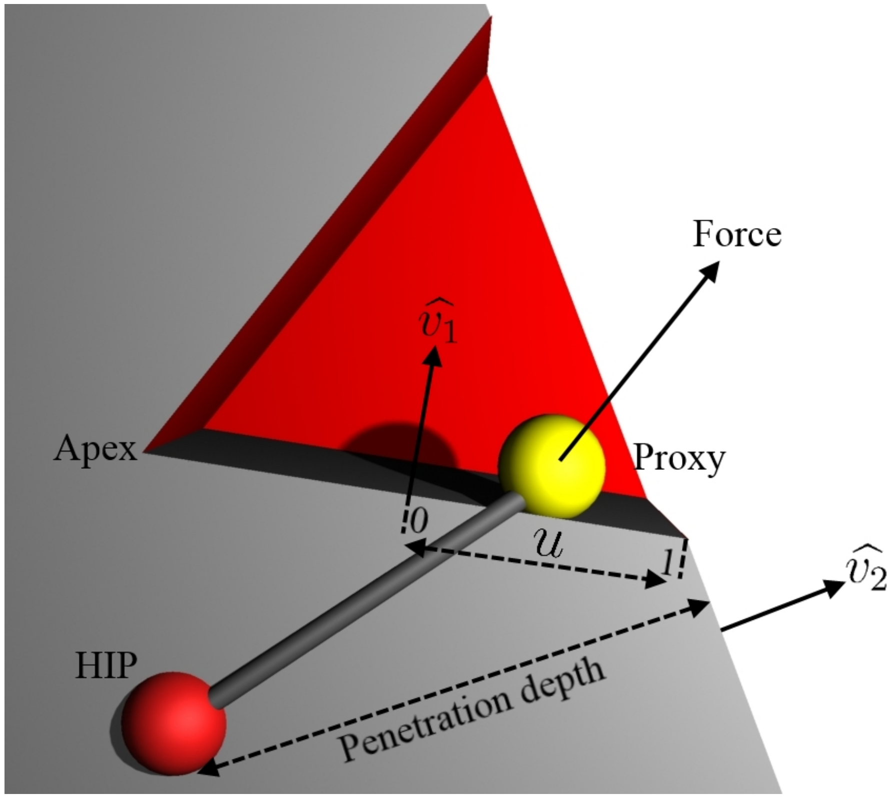

Figure 3.19: The SLERP of a deformable haptic cone with the virtual plane.

Consider two vectors $\widehat{v_1}$ and $\widehat{v_2}$, and find the angle between them:

$\omega = cos^{-1}(\widehat{v_1} \cdot \widehat{v_2})$

Given a parameter $u \in [0, 1]$, the slerp is:

$$slerp(u, \widehat{v_1}, \widehat{v_2}) = \widehat{v_1}\frac{sin((1-u)\omega)}{sin(\omega)} + \widehat{v_2}\frac{sin(u\omega)}{sin(\omega)} \quad (3.4.3)$$

Given that the proxy will slide towards the apex whilst deforming a cone, it is almost certain that the click will be performed at the target centre. The clamping at the apex means that it is very unlikely that the operator will slip off the target (a problem noted by Brewster et al. [Bre98]). Although there is a physical workload cost

Algorithm 3.1 The deformable cone implementation.

Require: The surface normal of the virtual plane ($\widehat{v_2}$)

```
while Haptic loop do
  if Not engaging with a deformable cone then
    Force = 0
    for all Deformable cones in the interface do
      if The HIP is interior to cone then
        Set engaging flag to true
        Position proxy on the surface
      end if
    end for
  else
    Constrain SCP to surface of the current cone
    Set the cone apex depth equal to the HIP depth
    Compute the Direction vector (Proxy - HIP) and normalise
    Compute the Penetration_Depth of the HIP to the surface of the virtual plane

    Displacement = Direction × Penetration_Depth
    F̂1 = Displacement × Stiffness
    Force = F̂1
    F̂2 = (Proxy - HIP) × Stiffness

    if F̂2_x > F̂1_x then
      Force_x = F̂2_x
    end if

    if F̂2_y > F̂1_y then
      Force_y = F̂2_y
    end if

    v̂1 = Force

    if The Proxy lies closer to an edge than to the cone apex then
      Compute the proxy distance from the edge u ∈ [0, 1]
      Direction = slerp(u, v̂1, v̂2)
      Force = Direction × ||Force||
    end if
  end if

  return Force
end while
```

(In the printed algorithm, Force, Direction, Displacement and (Proxy - HIP) carry an overbar denoting vectors.)

associated with deforming a cone, it is anticipated that the ability to navigate the interface with less intrusion from distracters will outweigh this. If the operator accidentally enters a deformable cone then they are provided with assistance when exiting it, in the form of the restoring spring force. As the operator climbs the cone wall it will begin to reform, which reduces the slant angle and makes it easier to exit. Video 9 demonstrates deformable cones in operation http://youtu.be/GiIDExUjv5c. Video 10 illustrates the effect of target distracters http://youtu.be/GCxWCYLFUN0.

3.4.6 Haptic virtual switch

To accurately position the cursor within a target and operate the device switch can be a challenge for many people with physical disabilities. Section 2.3.1 identified a number of difficulties that motion-impaired computer users encounter when clicking. The virtual switch proposed in this thesis is designed to simulate a push-button switch through haptic feedback. A virtual switch is placed around each icon within the interface and embedded into the virtual plane. The concept is based on existing assistive technologies, such as keyguards, that are designed to reduce unintentional key presses. The keyguard is a metal or plastic plate that is overlaid on top of the keyboard. The operator activates individual keys by poking through access holes.

One of the major difficulties that motion-impaired operators encounter is slipping off a target whilst clicking [Bre98] [TKM06]. It is likely that a flat surfaced switch would encounter this issue. The haptic cone proposed in Section 3.4.3 is an effective method of clamping the proxy to the target centre without imposing a force on the operator. Therefore, a pyramid shaped cone has been used for the surface of the virtual switch. The concept of the technique is presented in Figure 3.20.

Snap-action tactile switches are used in industry to provide decisive feedback. The simulation of a tactile switch will help provide appropriate haptic feedback to

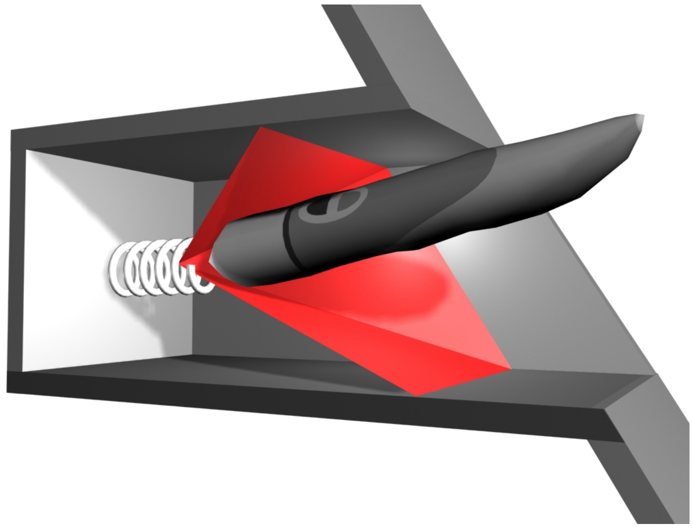

Figure 3.20: The concept of pressing a haptic virtual switch.

the operator. When engaging a virtual switch, the user will experience a restoring force until the spring is fully compressed, as shown in Figure 3.21(a). Once the spring is fully compressed, the force is disengaged until the switch reaches home, as shown in Figure 3.21(b). This provides the snap-action feedback of a tactile switch. The click is registered once the switch reaches home and an audio accompaniment confirms that the operation has been successful. Figure 3.21(c) shows the spring force that helps restore the operator to the surface of the virtual plane. To avoid switch bounce the release is only registered once the spring reaches a third of its restored displacement. This is accompanied by audio feedback to confirm the operation is complete.

In terms of the haptic rendering it is not necessary to translate the switch surface because the feedback of the restoring spring force can be computed based on the penetration depth of the HIP in relation to the cone apex. The initial contact force is calculated using the God-object algorithm but the z component is overridden when engaging a switch. The operator will still feel the sensation of pressing the switch but the cone will remain stationary in the haptic space. The main benefit of this approach is that there are no moving vertices and so the God-object algorithm will not suffer from pop-through. The cone surface can be translated in screen space to give the illusion of pressing the switch. The final implementation is presented in Algorithm 3.2.

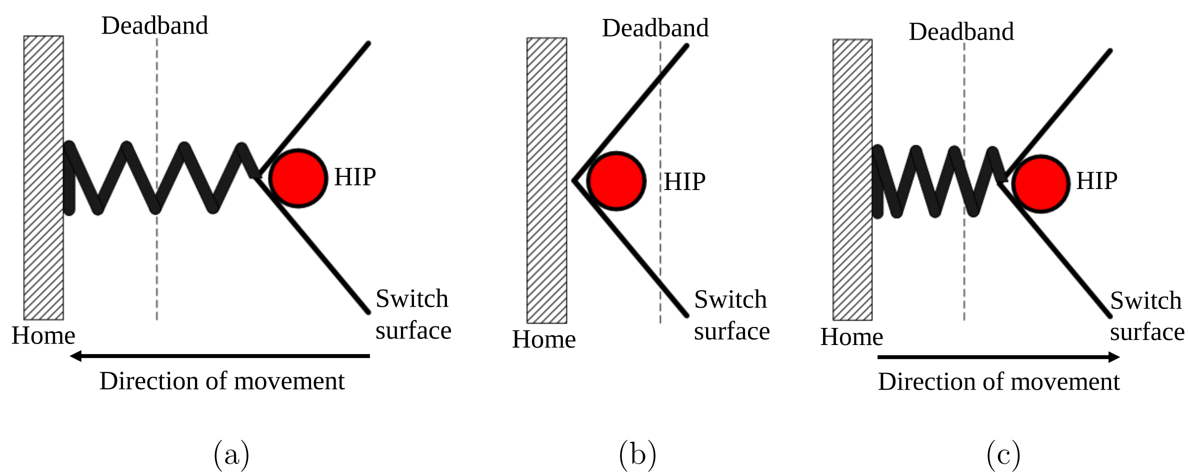

Figure 3.21: The operator presses the surface of the switch opposed by a spring force (a). When the HIP reaches the deadband, the force is disabled (b). Once the switch reaches home then the restoring spring force is engaged (c).

Trewin et al. state that accidental clicks are a major source of error for people with motion impairments [TKM06]. One of the advantages of the virtual switch is that it requires a conscious effort to operate and so it is less likely that a missed-click will occur accidentally. Clicking operations can only be performed on areas that contain virtual switches and so accidental presses of the device switch will be filtered out. When engaging a virtual switch, the proxy is guided towards the apex, which means that it is almost certain that the click will be performed at the target centre. It is very unlikely that the operator will slip off a target due to the effective clamping at the apex. Video 11 demonstrates virtual switches in operation http://youtu.be/R0_gPU8k43Y. Video 12 illustrates the effect of target distracters

Algorithm 3.2 The haptic virtual switch implementation.

Require: A virtual plane embedded with haptic cones

```
while Haptic loop do
  Calculate Force using Zilles and Salisbury's God-object algorithm [ZS95]
  if Not engaging with a virtual switch then
    for all Virtual switches in the interface do
      if The HIP is interior to a virtual switch then
        Set engaging flag to true
      end if
    end for
  else
    Translate the cone that is drawn in screen space to the depth of the HIP
    if The HIP reaches home then
      Set the down flag to true
      Play audio to confirm the switch is down
    end if
    if Not down then
      if The HIP has not reached the deadband (Figure 3.21(b)) then
        Compute the distance of the HIP from the apex
        Force_z = distance × stiffness
      else
        Force = 0 in the deadband
      end if
    else
      Compute the distance of the HIP from the apex
      Force_z = distance × stiffness
      if The HIP has reached 1/3 of the spring's displacement then
        Set the up flag to true
        Play audio to confirm the switch has been released
      end if
    end if
  end if
  return Force
end while
```

http://youtu.be/eNJU7jBdIfU.

3.4.7 Deformable virtual switch

The deformable virtual switch combines the virtual switch and the deformable cones discussed previously. If the operator wishes to select a target then they can press into the virtual plane at locations containing deformable switches, as shown in Figure 3.22(a). The deformation of the cone is included in the initial travel of the switch when the operator begins to compress the spring, as shown in Figure 3.22(b). Once the cone is fully deformed there is slightly further travel before the switch reaches home, as shown in Figure 3.22(c). The snap action then occurs and is processed in the same way as the virtual switch. When engaging the deformable switch, the proxy will slide towards the apex, which ensures that the click will be performed at the target centre. The clamping at the apex means that it is very unlikely that the operator will slip off the target. The most significant difference in the implementation of the deformable switches compared to the virtual switches is that Algorithm 3.1 is used to calculate the initial contact force rather than Zilles and Salisbury's God-object algorithm. The final implementation is presented in Algorithm 3.3. Video 13 demonstrates deformable switches in operation http://youtu.be/8yx2mQvStnU. Video 14 illustrates the effect of target distracters http://youtu.be/qGWvWwBCfuU.

3.4.8 Haptic workbox

Introduction

In recent years, typical desktop computer screen sizes and resolutions have increased significantly. The consequence of this is that a pointing device has to travel a much greater distance to navigate the whole of a computer screen. The additional workload can make a computer inaccessible for motion-impaired operators that suffer from fatigue or have a limited range of movement. Acceleration curves have been developed

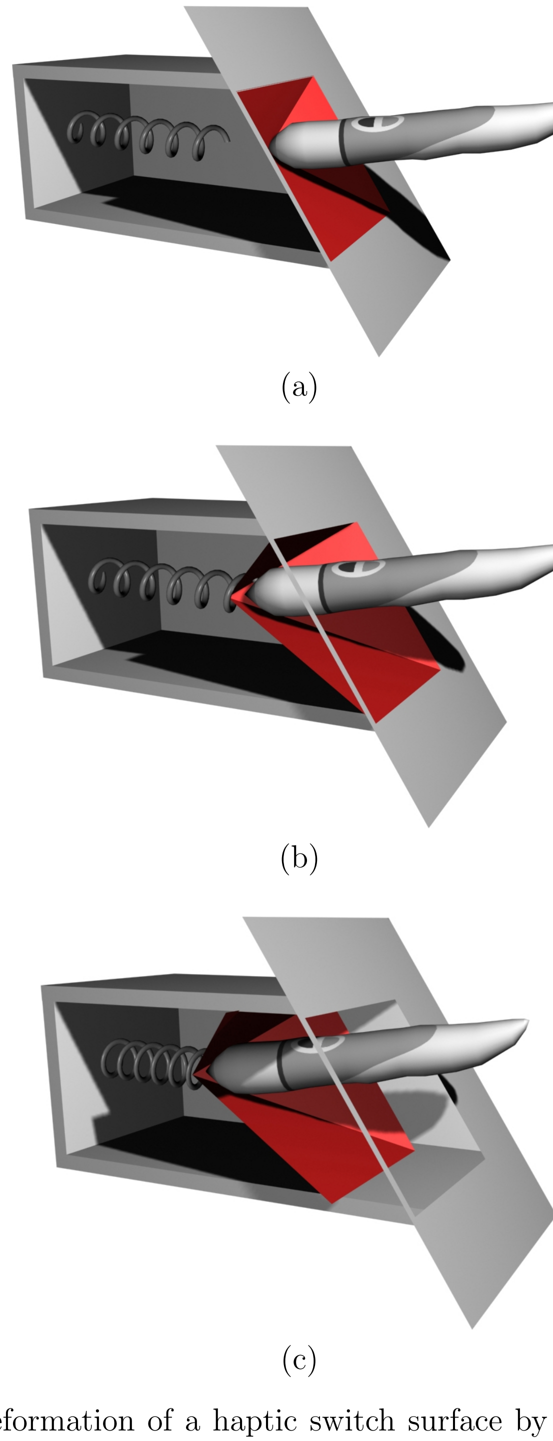

(a)

(b)

(c)

Figure 3.22: The deformation of a haptic switch surface by the virtual tool (a). A fully deformed haptic switch surface (b). Compressing the spring of a deformable virtual switch once the surface is fully deformed (c).

Algorithm 3.3 The deformable virtual switch implementation.

Require: A virtual plane embedded with deformable haptic cones

```
while Haptic loop do
  Calculate Force using Algorithm 3.1
  if Not engaging with a deformable virtual switch then
    for all Deformable virtual switches in the interface do
      if The HIP is interior to a deformable virtual switch then
        Set engaging flag to true
      end if
    end for
  else
    Translate the cone apex to the depth of the HIP
    if The HIP reaches home then
      Set the down flag to true
      Play audio to confirm the switch is down
    end if
    if Not down then
      if The HIP has not reached the deadband then
        Compute the distance of the HIP from the apex
        Force_z = distance × stiffness
      else
        Force = 0 in the deadband
      end if
    else
      Compute the distance of the HIP from the apex
      Force_z = distance × stiffness
      if The HIP has reached 1/3 of the spring's displacement then
        Set the up flag to true
        Play audio to confirm the switch has been released
      end if
    end if
  end if
  return Force
end while
```

to adjust the gain of the mouse when transitioning between coarse and fine navigation. However, haptic devices have a limited workspace and so predominantly use a direct mapping to the screen space, as depicted in Figure 3.23. The limitation of a direct mapping approach is that a higher gain is often required to allow the operator to navigate the whole of a computer screen. An increase in gain reduces the effective width of the targets and makes them more difficult to select.

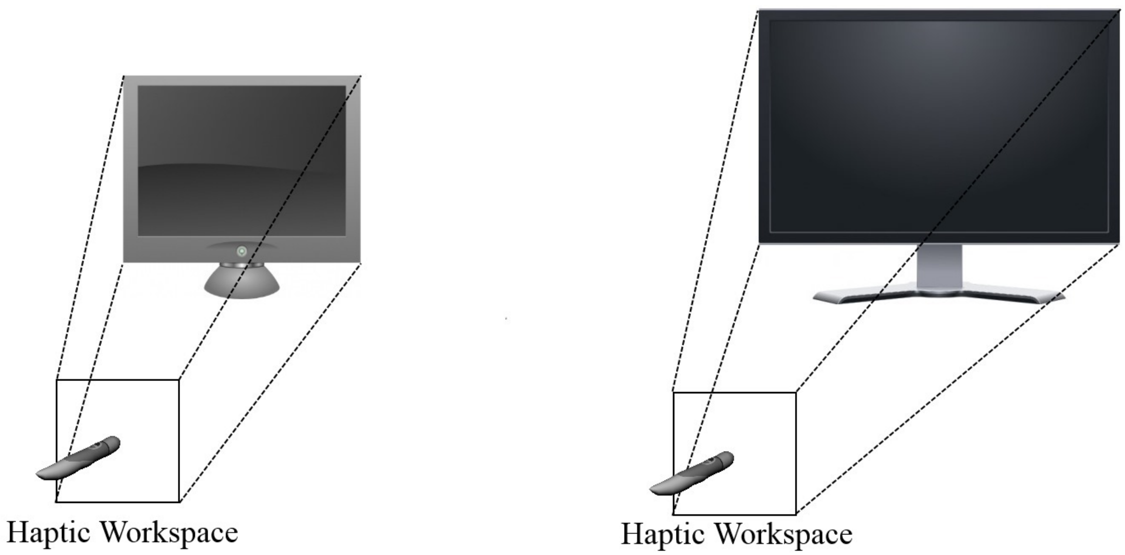

Figure 3.23: Direct mapping from the haptic workspace to screen space.

Rate control pointing devices, such as joysticks, allow the operator to navigate large areas of the screen with small input movements. This has the potential to benefit users that have a limited range of movement. However, a number of studies have shown that it is more difficult to accurately select small targets using rate control [Epp86] [CEB87] [MD96]. This is undesirable for motion-impaired operators because they often have more difficulty selecting a target than navigating towards it. [LHK+02b]. The workbox proposed in this thesis aims to utilise a rate/position hybrid system to combine the benefits of both input methods. Coarse navigation is rate controlled to allow the operator to navigate the whole of the computer screen and fine navigation is position controlled to allow accurate target selection. The workbox has the additional benefits of allowing small targets to be scaled up and can help reduce the effect of target distracters.

Implementation

The workbox can be considered as a parallelepiped workspace in which the user will interact. The coarse navigation of the cursor is rate controlled. This is achieved by pressing the proxy against the wall(s) of the workbox in the corresponding direction that the operator wishes the cursor to move. For example, pressing the proxy against the right hand wall will move the cursor to the right hand side of the screen, as shown in Figure 3.24. The user will be able to navigate the whole of the computer screen whilst only moving within the confinements of the workbox.

The rate of cursor movement is proportional to the force that the operator applies to the wall. Stocks et al. decided not to apply forces to the walls of their navigation cube so as to avoid confusion with other haptic interactions [SHL09]. However, during the preliminary study described in Section 4.2.3 the motion-impaired participants reported that the device felt too free without force feedback and the cursor would often overshoot the target region as a result. The restoring force of the walls provides essential stability to the hand for more accurate positioning of the cursor. The haptic cues discussed previously in this thesis can be integrated to assist target selection. These will be translated in the opposite direction to the cursor movement so that they are accessible within the workbox volume. All haptic cues will be disabled when the proxy is in contact with a wall so as to eliminate the effect of potential target distracters that may pass through the workbox. The rate control is temporarily disabled when the device switch is pressed to ensure that the cursor does not move if the operator accidentally presses against a wall when attempting to select a target.

The main difference between the workbox and the bubble approach proposed by

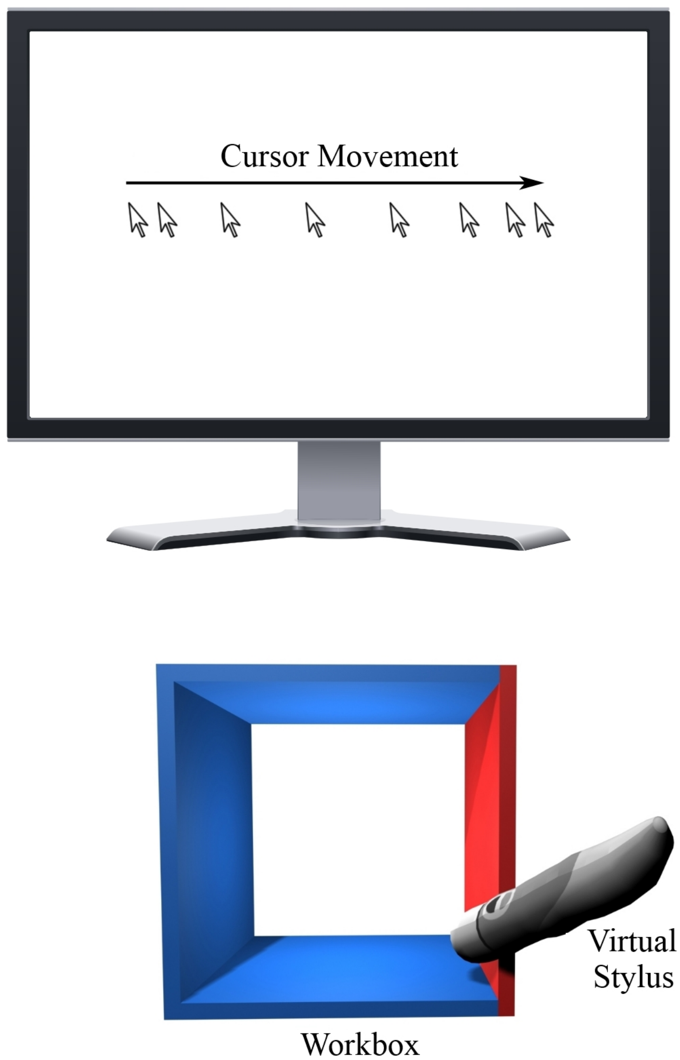

Figure 3.24: An example of how to navigate the cursor to the right hand side of the screen by pressing the tool against the corresponding wall of the workbox.

Dominjon et al. is the shape of the workspace [DLB+05]. The spherical bubble may be difficult to navigate for people with uncoordinated movements because they may wish to move the cursor in one axis and unintentionally engage the other axis at the same time. The separate walls of the workbox allow the user to navigate a single axis at a time. Additionally, Stocks et al. state that the spherical navigation volume in the bubble approach does not correspond well to the workspace of the haptic device and so they propose a navigation cube that is automatically scaled to fit the workspace [SHL09].

It is often desirable to scale up key features of an interface to assist target selection during the positional phase. For example, increasing the size of an interface can help people that have difficulty selecting small targets. This would not be permissible without the inclusion of rate control because the operator may not be able to reach the extremities of the screen with direct mapping alone. The black square surrounding the cursor in Figure 3.25 gives a representation of an area of the screen that has been magnified in the semi-transparent window overlaid on top of the workbox. The calculator in Figure 3.25 has been scaled by a factor of three in the haptic workspace. The repeated cursor within the magnified window maps directly to the tip of the virtual tool. Displaying the magnified window is not essential for user interaction but is useful for giving a visual representation of the haptic assistance and the scaling of key features within the interface.

Once the user has placed the black square around the desired target region then they can begin the target selection phase. The operator has direct positional control of the cursor within the workbox volume when the proxy is not in contact with a wall. This allows precise manipulations to be performed that would not be possible with rate control alone. During the positional phase the haptic cues surrounding the targets are re-enabled to aid target selection. The semi-transparent window allows

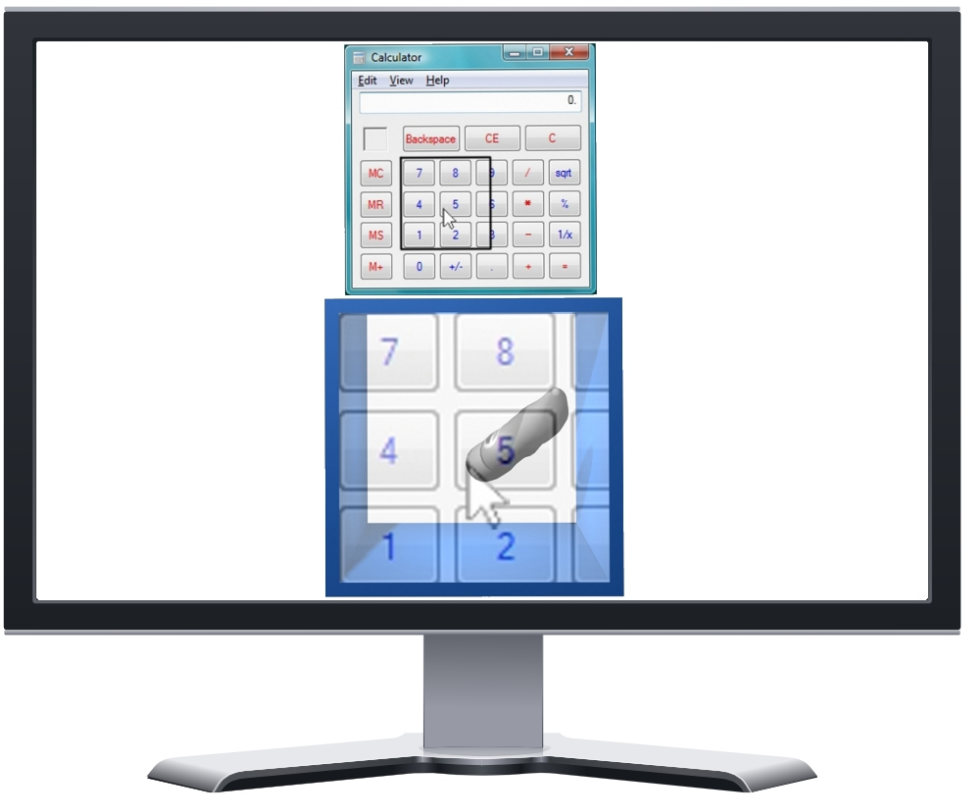

Figure 3.25: The black square surrounding the cursor indicates an area of the screen that will be magnified in the haptic workspace and screen space. The magnified semi-transparent window is overlaid on top of the workbox.

any visual cues that accompany the haptic assistance to be displayed in the OpenGL window behind.

User comfort is essential for motion-impaired operators especially when using a pointing device for long periods of time. A diagonal plane has been placed at the back of the workbox to provide a comfortable leaning position, as shown in Figure 3.26(a). The workbox has been positioned at the lower region of the y-axis so that all four walls are within reach when utilising the wrist rest supplied with the Phantom Omni, as shown in Figure 3.26(b). The size of the workbox, magnification level and rate gain can be adjusted to meet the needs of the individual. For example, a person with a very limited range of movement may require a smaller sized workbox to allow them to reach all four walls. A higher level of magnification may be useful for people that find it difficult to select small targets. The final implementation of the workbox is presented in Algorithm 3.4. Video 15 demonstrates the workbox in operation http://youtu.be/TiOe2UajXds.

Algorithm 3.4 The haptic workbox implementation.

```
Require: A parallelepiped workbox
Require: A magnification factor
Require: Haptic cues e.g. gravity wells

  Scale the haptic cues by the magnification factor.
  while Haptic loop do
    if The HIP lies within the workbox then
      Enable the haptic cues.
      Cursor position = Proxy position.
    else
      Render the restoring force for the walls.
      Disable the haptic cues.
      if The device switch is not down then
        Translate the cursor and black square proportional to the force.
        Set the position of the magnifier.
        Translate the haptic cues in the opposite direction.
      end if
    end if
  end while
```

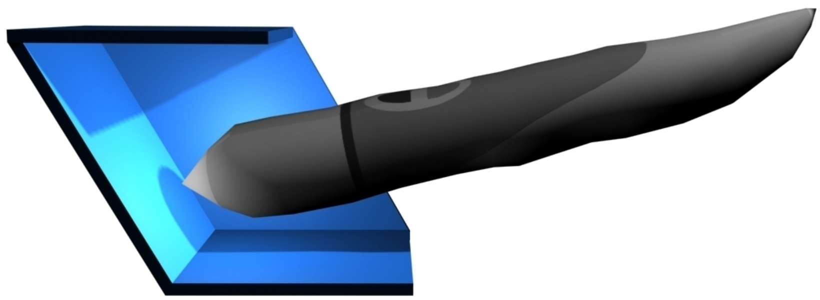

(a)

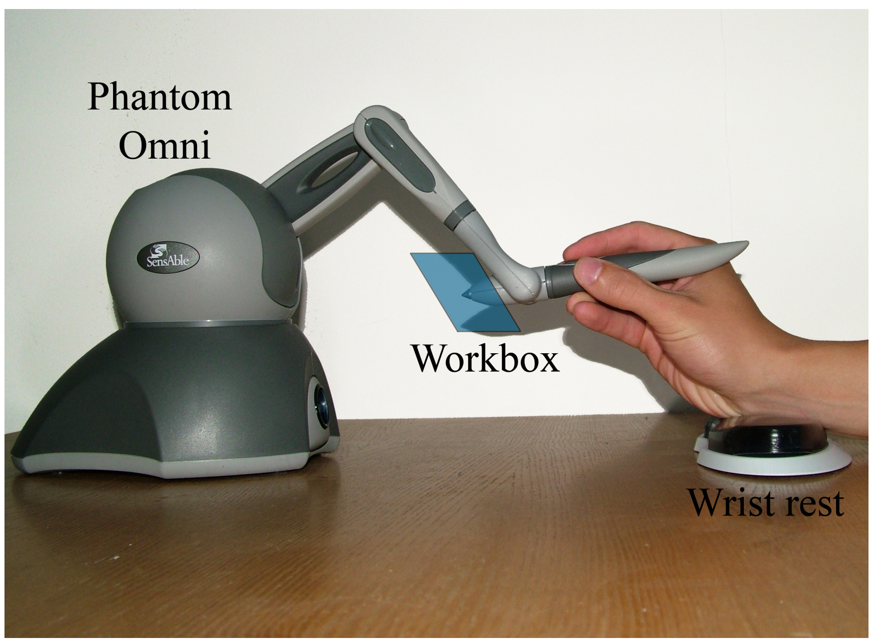

(b)

Figure 3.26: A side view of the parallelepiped workbox with the virtual stylus (a). A scale view of the workbox and its position within the physical workspace of the Phantom Omni (b).

3.5 Summary

This chapter has presented the implementation of the haptic assistance investigated throughout the remainder of the thesis. The literature highlighted that the majority of difficulties regarding point-and-click tasks lie primarily in the final stages of target acquisition. Therefore, five novel haptic conditions are proposed to improve the selection of small targets including: haptic cones, V-shaped funnels, deformable cones, virtual switches and deformable switches. These techniques are designed to address many of the concerns that were highlighted in Chapter 2. A new haptic rendering algorithm has been developed to permit the implementation of the deformable assistance. Finally, a rate/position hybrid system called the haptic workbox has been proposed to permit rapid cursor navigation whilst allowing accurate target selection. The workbox is also designed to reduce the physical workload and limit the effect of target distracters. The following chapter describes the experimental setup that has been used to evaluate the haptic assistance presented in this chapter.

Chapter 4

Experimental setup

4.1 Introduction

The majority of human-computer interaction is performed using a pointing device to navigate the on-screen cursor [DJ06]. Haptic feedback is designed to improve human motor performance in a virtual environment. This chapter describes the experimental setup that has been used to evaluate the effectiveness of the haptic assistance presented in this thesis.

4.2 Point-and-click tasks

The following section describes the implementation of the point-and-click tasks that have been conducted in this study as part of the evaluation process. The most common point-and-click tasks in the literature are variants of the ISO 9241-9 standard for pointing device evaluation. The implementation of this multidirectional experiment is described in Section 4.2.1. The ISO 9241-9 task has a number of shortcomings that limit its effectiveness for evaluating haptic assistance. As a result, an additional experiment is proposed in Section 4.2.2 that will more vigorously test haptic assistance in a real world GUI.

The experimental tasks were explained and demonstrated to the participants prior to the study being conducted. In accordance with Fitts' law participants were instructed to work as quickly as possible whilst maintaining a high level of accuracy. The user group were informed that they could rest at any time between trials. The on-screen cursor was projected on top of the x-y position of the proxy so that the visual feedback corresponds to that experienced by the user through the haptic device.

When performing empirical studies with motion-impaired computer users, practical limitations can restrict the application of detailed statistical analysis. The main limitations involve the increased heterogeneity of motion-impaired users compared to able-bodied ones and the small sample size. It is necessary for researchers to run the trials on a long-term basis, to develop a working relationship with the users and to keep experimental conditions constant [HKL+01]. Repeated measures designs are generally employed because of the small number of users available.

4.2.1 ISO 9241-9 experiment

The ISO 9241-9 standard for computer pointing devices proposes an evaluation of performance and comfort [PDE99]. Graphical user interfaces (GUIs) have evolved and matured significantly since their introduction with the Apple Macintosh in 1984. The key feature of modern GUIs is the ability for users to interact with simple point-and-click operations. The ISO 9241-9 establishes uniform guidelines and testing procedures for evaluating computer pointing devices. The metric for comparison is throughput, in bits per second (bits/s), which includes both the speed and accuracy of the users (See Section 2.7.1). The ISO 9241-9 was in Draft International Standard form in 1998 and became an International Standard in 2000. The primary motivation was to influence the design of computer pointing devices to accommodate the user's biomechanical capabilities and limitations, allow adequate safety and comfort, and prevent injury. Secondarily, the standards establish uniform guidelines and testing procedures for evaluating computer pointing devices produced by different manufacturers. Compliance is demonstrated through the testing of user performance and comfort to show that a particular device meets ergonomic requirements.

The multidirectional point-and-click task described in the ISO 9241-9 standard consists of evenly spaced targets arranged in a circular layout. The number of targets and their shape can be adjusted to the desired specification. The task requires the participant to first click on the top target, then on the target directly opposite, then the next target clockwise in the sequence and so on around the circular layout. This paradigm has the advantage of controlling for the effect of direction. Figure 4.1 shows the layout and target sequence that was conducted in the experiments presented in Section 5.4. The radius of the circle was 50mm in device displacement. Data collection begins once the first target has been selected and continues until the task is completed. The current target in the sequence is highlighted in red and only progresses to the next target once acquired. A click is only registered if the proxy is in contact with the virtual plane. This ensures that the operator has used the assistance and its benefits can be recorded.

4.2.2 On-screen keyboard (OSK) experiment

The ISO 9241-9 task has limitations when analysing haptic assistance because it does not take into consideration the effect of target distracters [ADL11]. In most GUIs the toolbar buttons are arranged in rows or columns and so the evenly spaced, circular layout is unrealistic. A number of the studies discussed in Section 2.6.3 have shown that the effectiveness of haptic assistance can be significantly reduced when distracters are in close proximity of a target.

Many assistive technologies for people with physical disabilities either utilise the keyboard or emulate its functionality. In this study the Windows on-screen-keyboard

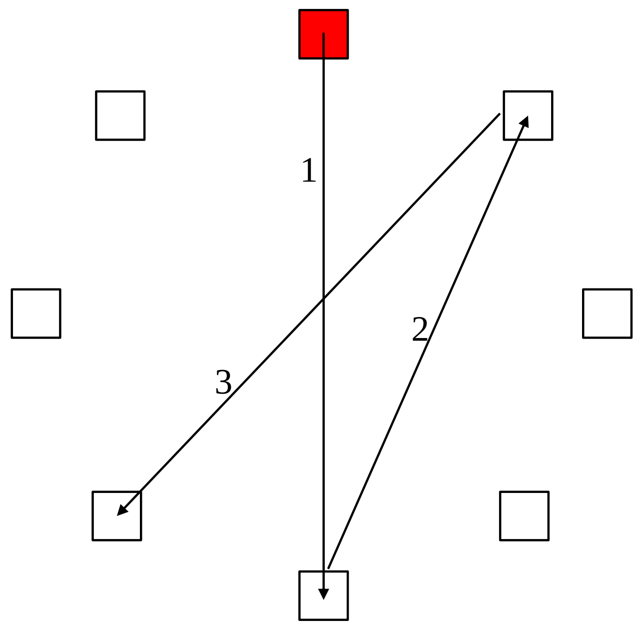

Figure 4.1: The layout and target sequence for the ISO 9241-9 multidirectional task.

(OSK) has been chosen as the primary interface to evaluate the haptic assistive techniques. The densely populated GUI will provide a more realistic and extensive evaluation of the effects of target distracters. The experimental task requires the participant to perform fifty successful selections to produce a predefined sentence using the OSK. The target key is highlighted in red to remind the participant which character is next in the sequence. The same sentence was used throughout the study to ensure that the index of difficulty did not change. This enables post-test comparisons to be made between haptic techniques for measures such as movement time and the experiment distance travelled. The structure of the sentence was: "WE ARE USING THE PHANTOM OMNI AND HAPTICS TO TYPE.". This sentence was chosen because it has a variation of direction changes and distance between letters. The participants were familiar with the structure of the sentence from the practice sessions. This was not an issue because in real world applications the user would know which letter they are going to type or the icon they are going to select.

The dimensions of the character keys were 7.1mm × 7.1mm in device displacement and the spacebar dimensions were 45.9mm × 7.1mm. A time limit was not imposed on the experiment and so the duration was dependent on the individual's ability to complete the task. Data collection begins once the first target is selected and continues until the sentence is completed. Any selections of surrounding keys were recorded for the cursor analysis but ignored in the textbox sentence. i.e. the operator was not required to delete undesired key selections. A click is only registered if the proxy is in contact with the virtual plane. This ensures that the operator has used the assistance and its characteristics can be recorded.

Visual feedback is provided to the operator through an OpenGL window to help support interaction with haptic cues. The opacity of the OSK window is adjusted to make it semi-transparent, which allows the features to be seen behind, as shown in Figure 4.2. Additional information on the window transparency and depth cues is provided in Section 4.5. The position of the keys is obtained using the interface feature extractor discussed in Appendix B.2.

4.2.3 Experimental preparation

Although the ISO 9241-9 standard does not discuss learning effects, this obviously must be considered when designing and evaluating performance data. To ensure that the participants were familiar with using the Phantom Omni a number of experimental tasks were performed. Over a twelve week period the participants were encouraged to use the Phantom Omni in gaming and simulation environments. Typical tasks involved colour pairing, target shooting, haptic archery, haptic basketball and a haptic xylophone. Figure 4.3 illustrates a number of screenshots from these applications. The tasks were designed to be engaging and emphasise the 3DOF capabilities of the Phantom Omni.

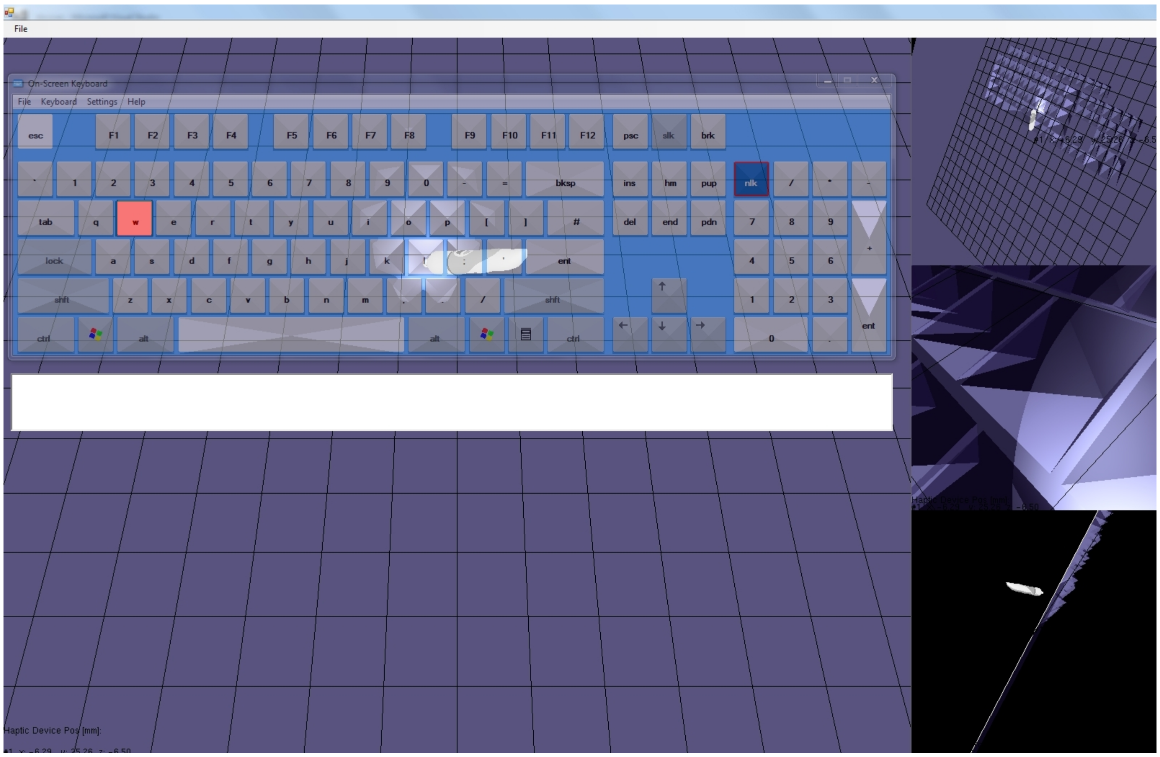

Figure 4.2: The semi-transparent Windows on-screen keyboard interface with the OpenGL window behind.

Another session ran in parallel where participants could familiarise themselves with the experimental task and each type of haptic assistance. In this study the participants performed the on-screen keyboard point-and-click task described in Section 4.2.2 for ten blocks of 50 trials each before the data was recorded for analysis. It was hypothesized, based on other similar experiments, that they would have achieved a criterion level of practice by block ten [DM94] [DKM99] [MKS01]. Essentially, no significant improvement in performance would be shown in the final blocks. The sessions were typically shared between the user group over a weekly two hour session for the twelve weeks. After this time the participants were familiar with the Phantom Omni and the experimental tasks. The analysis of the 3 repetitions for each haptic condition in Experiment 1 confirms that there were no learning effects or issues

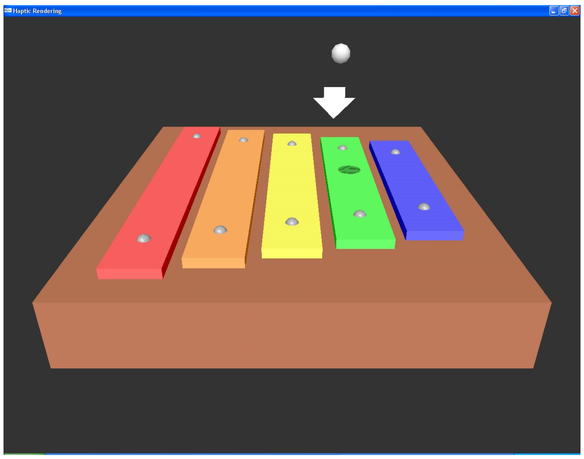

(a) Haptic xylophone. The operator performs a musical tune by hitting the key that the arrow points to.

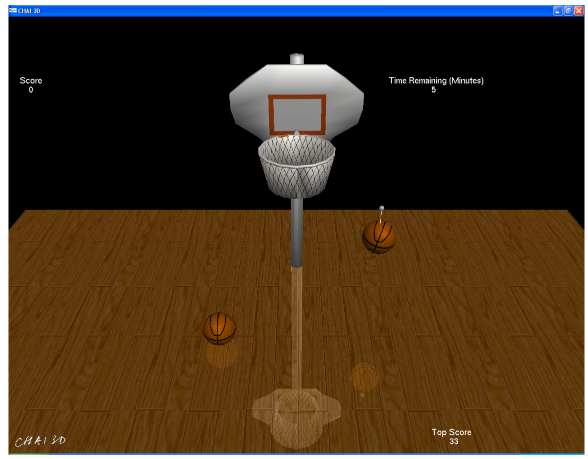

(b) Haptic basketball. The operator has to pick up the basketballs and place as many through the hoop as possible in 5 minutes.

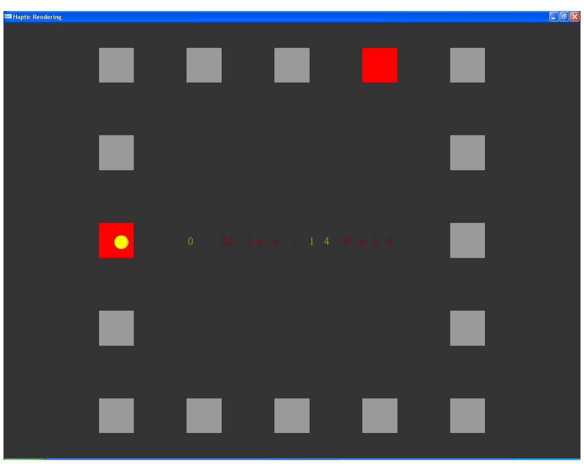

(c) Colour pairing. The operator has to match the coloured pairs in the fastest time possible. Haptic cues can be added to the squares if required.

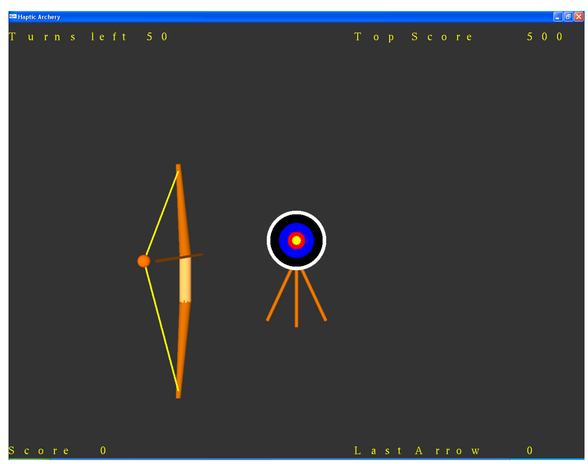

(d) Haptic archery. The operator clicks, pulls back the bow and then releases the device switch to fire the arrow at the target.

Figure 4.3: Examples of the games and puzzles that were used to familiarise the participants with the Phantom Omni.

regarding fatigue.

4.2.4 Hardware

The experiments conducted in this thesis were performed using a Dell Precision M4500 laptop computer. The technical specifications are presented in Table 4.1.

Operating System	64-bit Windows 7
Processor	Intel® Core^{TM} i5-540M Processor (3M Cache, 2.53 GHz)
Graphics Card	NVIDIA® Quadro FX 1800M Graphics with 1GB2 dedicated memory
Screen	39.6cm (15.6”) FHD LED Back-Lit Display (1920×1080) Resolution
RAM	8GB

Table 4.1: The technical specifications of the equipment used in the study.

4.2.5 Data collection and playback

The raw data from each experiment is written to file to allow future analysis to be performed. The relevant information is collected on the interface layout and the position and state of the pointing device. This allows the researcher to create new cursor analysis techniques to evaluate new areas of interest. The appropriate data needs to be collected to perform the cursor analysis techniques discussed in Sections 2.7 and 3.2. However, writing data to file is a slow process and so it is important to ensure that there is not a latency that could potentially affect the haptic and graphical interactions. The data collected in each log file is listed below:

The following data was recorded on initialisation of the experiment.

- Assistance type: (Gravity well, Haptic cone, Deformable switch, etc.)
- Participant: (ID)
- Task type: (ISO / OSK)
- A list of interface targets: (Rectangles)

The following data was recorded for each pointing operation.

- Current target position: (Rectangle)
- The current position of the proxy: (x, y, z)
- Button down position: (x, y, z)
- Button up position: (x, y, z)
- Proxy surface contact: (Boolean)

Measures, such as device velocity, movement time and the experiment distance travelled, are calculated retrospectively based on the change in position, the number of samples and the sampling rate. A file reader has been created that analyses the raw data using the cursor analysis techniques and exports a table of results. A playback feature allows the researcher to observe the cursor trace from previous experiments and identify the benefits or limitations of a certain haptic condition. Figure 4.4 shows an example of playing back an unassisted experiment. The file reader can be used to replay two files simultaneously. This allows the researcher to visually compare two different assistive conditions at the same time. The measures for each click are displayed in the lower right panel during playback.

The results from the cursor analysis will help determine if the newly proposed haptic techniques outperform the traditional assistance. The evaluation is based on reducing error rates, improving targeting times and eliminating or reducing the shortcomings discussed in Chapter 2. The following section describes the participants that took part in the study.

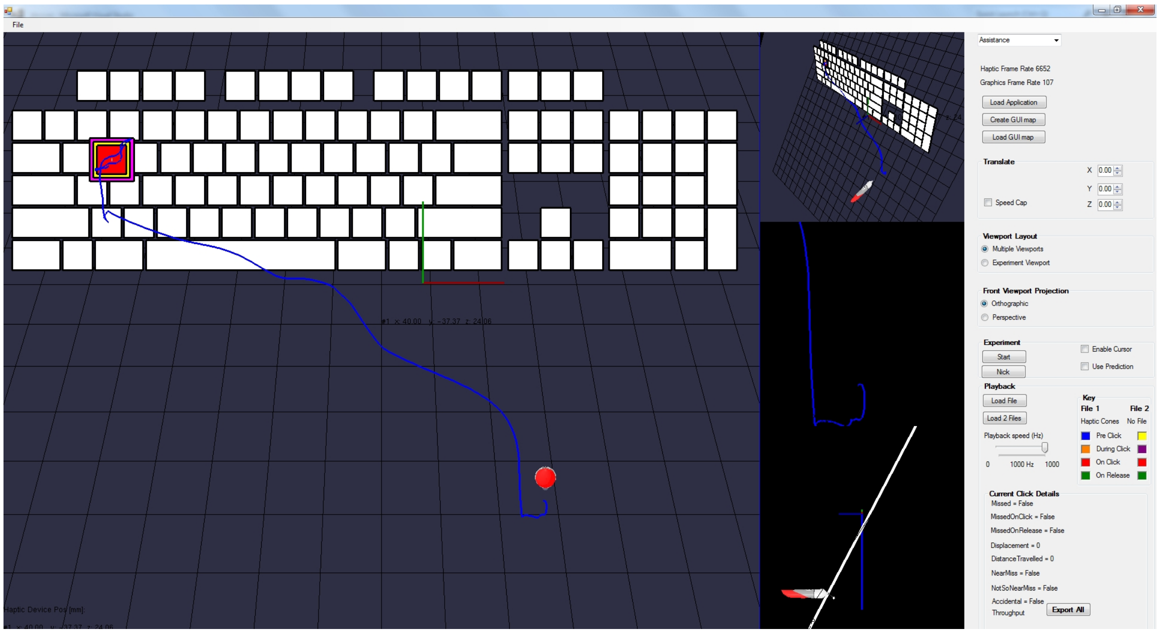

Figure 4.4: The file reader replaying an unassisted OSK experiment. The bottom right panel provides the metrics of the current clicking operation.

4.3 Participants

In this study seven volunteer participants (3 women and 4 men) with varying degrees of motion impairment were recruited from the Norfolk and Norwich Scope Association (NANSA). They were aged between 35 and 58 with average and median ages of 45 and 49 respectively. The users were affected by cerebral palsy (6) and spina bifida (1). The group represented a wide range of capabilities, exhibiting symptoms including tremor, coordination difficulties, stiffness, weakness and reduced dexterity in the dominant hand and arm. All the participants have had more than 5 years of computer experience and use a computer several times a week. The participants were both selected and screened for the experimental trials. Selection was managed to equally distribute the range of disabilities in the sample from high to medium and low impairments. Screening was also managed to exclude individuals with significant sensory or cognitive impairments that may have interfered with their capability to perform the tasks. People with motion impairments can experience changes in their capability over time. They become fatigued easily, despite extremely high motivation and sometimes cannot complete trials or conditions. It is necessary for researchers to run the trials on a long-term basis and to develop a working relationship with the users so that experimental conditions remain constant [HKL+01]. Repeated measures designs have been employed because of the small number of participants and their limited availability.

Ethical approval was obtained from the appropriate NANSA committee. Informed consent was obtained from each participant prior to the commencement of data collection. A key-worker was present at all times during the experimental sessions to ensure the wellbeing of the user group. A brief summary of the participants' physical background has been provided in Table 4.2. Please note that as a result of patient confidentiality, access to medical records and assessments is not permitted.

4.4 Device comfort

When operating a pointing device it must be comfortable to use for long periods of time. This is particularly important for people with motion impairments. The following subsections discuss the methods that have been employed to ensure user comfort when operating the Phantom Omni.

4.4.1 Positioning the Phantom Omni

The correct placement of the Phantom Omni in relation to the operator will vary from one person to another. It is essential that the device is positioned so that the user does not experience any strain on the wrist or forearm when working within the haptic workspace. In general the operator will be most comfortable when they can

Gender	Age	Disability	Details	Legend
Female	35	Cerebral Palsy	Can walk unaided, has speech difficulties and communicates through a communication aid. Principal impairment is tremor which makes finer movements difficult to perform.	□
Male	49	Cerebral Palsy	Can walk unaided, has poor co-ordination and finds it difficult to perform finer motor control.	△
Female	49	Spina Bifida	Is an electric wheelchair user, has very good fine motor control but often takes a long time to complete the tasks.	●
Male	58	Cerebral Palsy	Is an electric wheelchair user. Principle impairments are muscle stiffness, spasm, limited co-ordination and speech difficulties. The participant has difficulties locating the device switch.	×
Female	52	Cerebral Palsy	Is a manual wheelchair user. Movements can be quite slow but are reasonably well controlled. As a result, error rates are low but the task can take longer to complete.	○
Male	38	Cerebral Palsy	Can walk unaided. Principal impairments are tremor and spasm which makes finer movements difficult to perform.	+
Male	35	Cerebral Palsy	Is an electric wheelchair user. Has a very limited range of movement. Finds it difficult to grasp the stylus and cannot operate the device switch.	⋆

Table 4.2: A brief summary of the background of the six participants within the study.

rest their elbow on the desk and support their wrist using the supplied wrist rest, as shown in Figure 4.6. During the trials the participants were given the opportunity to position the device according to their own preference. The researcher ensured that the haptic workspace was central to the operator's range of movement so that the whole screen was accessible. Each participant gripped the stylus in accordance to their own personal preference. The height adjustable table shown in Figure 4.5 is a useful tool for positioning the device correctly for wheelchair users. The device guidelines advise regular breaks so that the user can stretch their hands, wrists and elbows.

Figure 4.5: A wheelchair user operating the Phantom Omni on a height adjustable table.

4.4.2 Haptic virtual plane

During the preliminary study described in Section 4.2.3 some participants experienced arm ache due to not having a surface to rest the stylus against. The response to this problem was to introduce a haptic virtual plane that the operator may lean against for additional support. An example of this is shown in Figure 4.6. The slant angle of the virtual plane provides a comfortable leaning position, similar to a graphical designer's drawing board. The technique is especially useful when working at the upper extremities of the workspace (where the wrist rest has less effect) because the user can place their elbow on the desk and rest against the virtual plane. A virtual ceiling is positioned in parallel to ensure that the proxy is always the same perpendicular distance away from the virtual plane regardless of the device's y-position. Many of the haptic assistive techniques presented in Chapter 3 are applied to the surface of the virtual plane or embedded in it. The following section describes the depth cues that have been implemented to improve depth perception when interacting with the virtual plane and other haptic assistance.

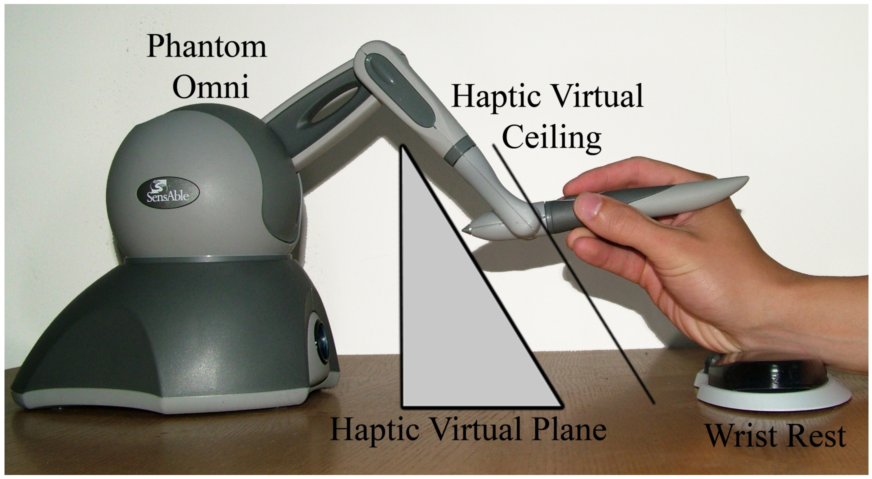

Figure 4.6: The concept of a haptic virtual plane.

4.5 Depth cues

Everyday visual perception involves interacting with a three-dimensional world, however, interaction with a computer is typically based on a two-dimensional display surface. When interacting with a virtual environment it is important that the visual information is sufficient so that the operator can perceive the depth of the tool in relation to the haptic cues. A number of depth perception techniques have been implemented to improve this aspect of interaction. The combined depth cues can be observed within the interface presented in Figure 4.2.

4.5.1 Tool spotlight

Initially the scene is rendered with a single light source located at the viewer's position. This light illuminates the surfaces with only ambient and diffuse lighting to provide three-dimensional definition but without a specular highlight. A second light source is applied to the end of the virtual tool in the form of a spotlight. This light source contains ambient, diffuse and specular components. The light is always directed along the Vector (0,0,1) to ensure that the spot size is proportional to the proxy's distance

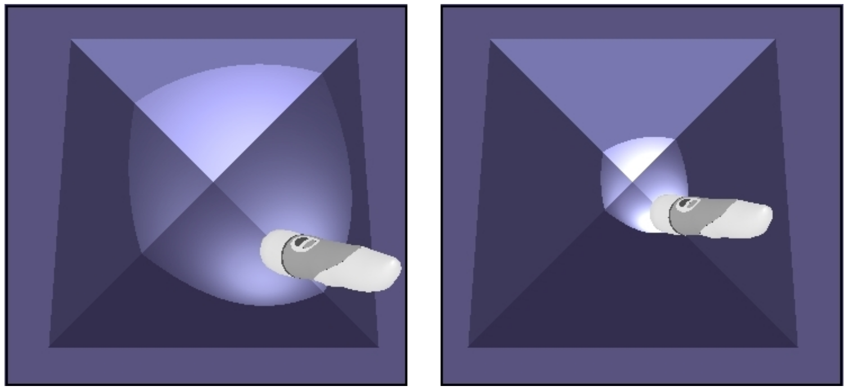

(a) (b)

Figure 4.7: The tool spotlight illuminating the surface of a haptic cone at a distance (a). The tool spotlight illuminating a haptic cone close to the surface (b).

away from a surface. This will help the operator perceive the depth of the tool in relation to the virtual plane and haptic cues. Figure 4.7 shows an example of the spotlight illuminating a haptic cone at different depths. Per-pixel lighting has been employed using OpenGL and the OpenGL Shading Language (GLSL).

4.5.2 Linear perspective and texture gradient

Linear perspective has been utilised to improve depth perception by overlaying a "texture gradient" grid on top of the virtual plane in a perspective viewport. The parallel lines facing the observer converge as they move away towards the top of the screen. The more the lines converge, the farther away they appear. Figure 4.8 illustrates how this approach can improve depth perception.

4.5.3 Virtual tool colour

When the operator makes contact with a surface, the virtual tool changes colour from grey to red, as shown in Figure 4.9. This provides confirmation to the user that they have reached the desired surface.

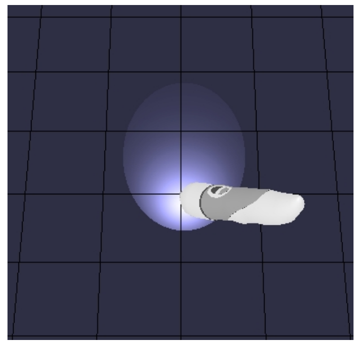

Figure 4.8: A "texture gradient" grid is placed over the virtual plane. The convergence of the lines helps to give a better perception of depth on the surface.

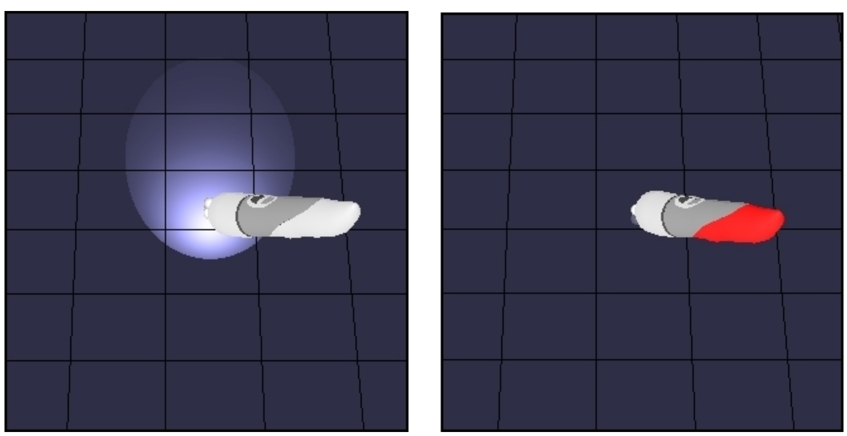

(a) (b)

Figure 4.9: When the proxy makes contact with a surface, the virtual tool changes colour from grey (a) to red (b). This provides confirmation to the user that they have reached the destination.

4.5.4 Window transparency

One of the difficulties when integrating a 3DOF haptic device with a two-dimensional graphical user interface (GUI) is providing visual depth information. The depth perception techniques that have been discussed previously are designed to improve this aspect of interaction. However, the GUI window and OpenGL window are separate, which means that the depth information cannot be conveyed directly. To overcome this issue the opacity of the GUI window can be adjusted to make it semi-transparent. The semi-transparent GUI window is then positioned on top of the main viewport in the OpenGL window. This allows the depth information to be perceived through the GUI. An example is provided in Figure 4.10 where the haptic cones can be seen through the on-screen keyboard.

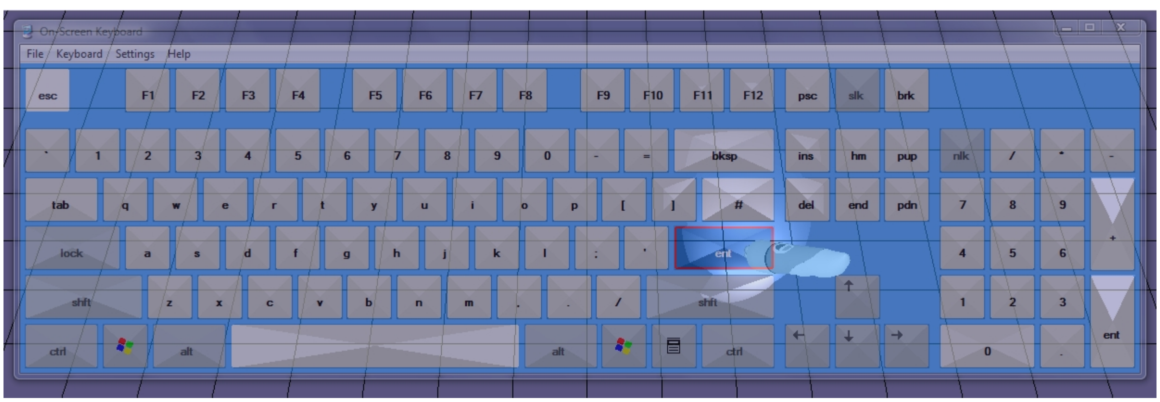

Figure 4.10: The semi-transparent on-screen keyboard is placed on top of the main OpenGL window. The haptic cones and depth information can be perceived through the GUI.

4.5.5 Multiple viewports

In recent years computer screens have increased significantly in size and resolution. This allows more information to be displayed on the screen at any given time. Therefore, multiple viewports can be introduced without significantly reducing the size of the main viewing area. Four viewports have been provided to allow simultaneous visualisation of a variety of view types and angles. This is designed to enhance the user's understanding of their position and orientation within the three-dimensional workspace. The multiple viewports are also useful to the researcher when making observations during experimental tasks. Figure 4.11 shows an example of the multiple viewports that are available to the user.

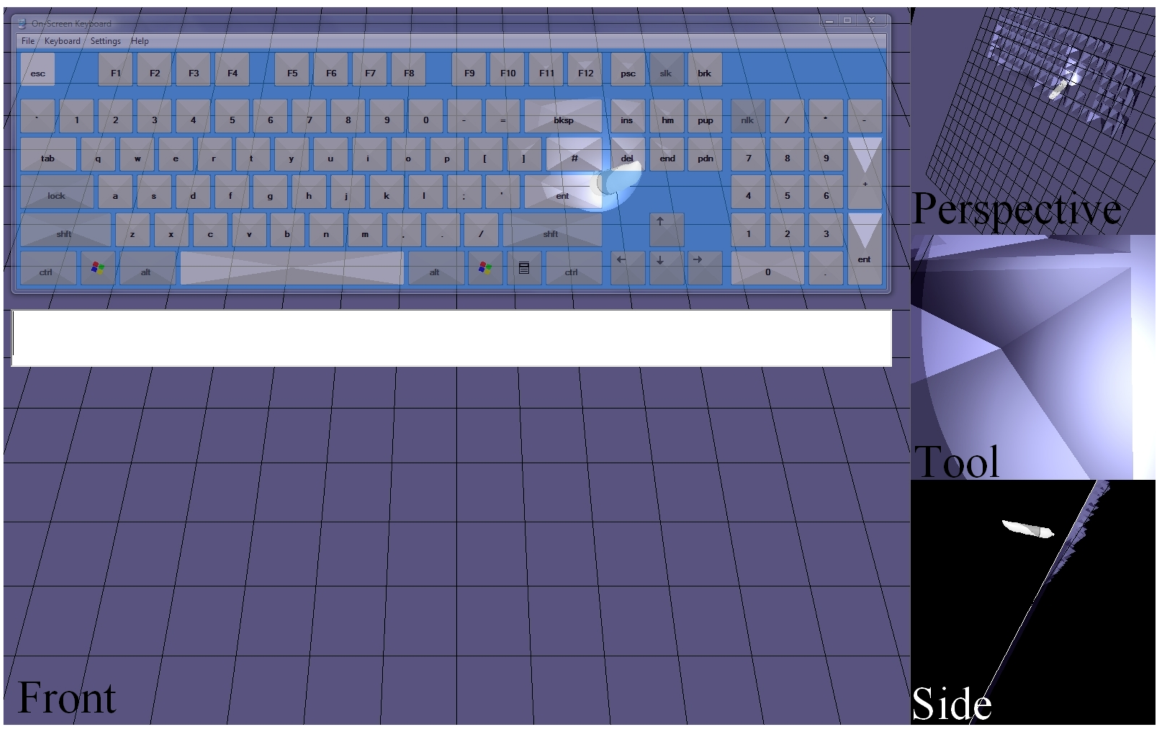

Figure 4.11: The four viewports give the operator a better understanding of their position within the haptic workspace. The multiple views are useful to the researcher when taking observations during experimental tasks.

Chapter 5

Results

5.1 Introduction

The results of the study are presented in this chapter, including data analysis and descriptive statistics. The results chapter is divided into four sections, each of which is devoted to a particular evaluation of haptic assistance. The structure is as follows:

- Experiment 1 investigates the performance benefits of haptic assistance that does not require force calibration compared to traditional techniques in a static configuration. Haptic cones and V-shaped funnels are compared against gravity wells and high-friction targets. The focus of the analysis is on error rates and targeting times.

- Experiment 2 explores the benefits of non-intrusive haptic assistance that the operator can choose to use or ignore. The study concentrates on the evaluation of deformable haptic cones and deformable virtual switches. The analysis is focused on interaction rates and the effect that target distracters have on user performance in a densely populated interface.

- Experiment 3 uses the ISO 9241-9 task to investigate how target size and shape effect the performance of the deformable haptic assistance. The study is designed to show that the techniques are generalisable for real-world GUIs that contain different target sizes and shape.

- Experiment 4 investigates the benefits of the haptic workbox in terms of improving computer access for people with severe motion impairments.

The results from these studies will help determine if the newly proposed haptic techniques described in Chapter 3 overcome the three major shortcomings of traditional haptic assistance that were identified in Chapter 2. The experiments were performed after the twelve weeks of practice sessions described in Section 4.2.3. Each experiment was observed by an assistant to ensure that it had been completed without any complications. For each cursor measure discussed in this Chapter (apart from throughput and experiment time spent on the virtual plane) a reduction in magnitude is desirable and will signify an improvement. All statistical analyses were performed in GraphPad PRISM version 6.02 (GraphPad, CA, USA). Unfortunately, subjective measures were unobtainable in this study due to communication difficulties of some of the participants. Please note that the participant denoted by $\star$ was only included in Experiment 4 part 3 because they were unable to accomplish the other tasks.

5.2 Experiment 1

5.2.1 Experimental procedure

Experiment 1 is designed to provide an insight into the effectiveness of the newly proposed haptic assistance in comparison to the traditional techniques. This study concentrates on the performance benefits of haptic cones and V-shaped funnels compared against gravity wells and high-friction targets. The aim is to determine if techniques that do not require force calibration outperform traditional haptic assistance in a static configuration. The force levels of gravity wells and high-friction targets were chosen to provide a suitable level of assistance for target acquisition whilst permitting the user to exit haptic cues without opposition from excessive forces. In this instance the spring constant that provides the gravity well clamping at the target centre was chosen to be 0.4. The coefficient of friction for high-friction targets was chosen to be 0.4 for the static and dynamic components respectively.

The six participants were asked to perform the on-screen keyboard (OSK) task described in Section 4.2.2. Each operator was required to create a predefined sentence within a textbox using the OSK. The active target was always highlighted in red and only progressed to the next key once acquired. A sequence comprised of 50 successful target selections and data collection began after the first target had been selected. Data collection was continuous within a sequence and breaks of a minimum of 10 minutes were taken between trials. In accordance with Fitts' law participants were instructed to work as quickly as possible whilst maintaining a high level of accuracy. Due to only having a small sample size the experiments were repeated three times for each haptic condition. The order of presentation of the haptic techniques was randomised for each person. For the unassisted condition the force feedback provided by the haptic cues was turned off, meaning the Phantom Omni operated as an "ordinary" pointing device. Sessions were two hours long and shared between the user group. The cursor movements of each participant were recorded using the data recorder discussed in Section 4.2.5 and analysed according to the measures described previously in Sections 2.7 and 3.2. The main emphasis of the cursor analysis is on error rates and targeting times. For each cursor measure discussed in this section (apart from throughput) a reduction in magnitude is desirable and will signify an improvement.

The results for each participant and haptic condition have been presented using interleaved scatter plots. The data has been described using the mean of the 3 repetitions and standard error of the mean (SEM). The participants' physical background and corresponding legend can be found in Table 4.2. To test the null hypothesis that there were no differences between haptic conditions for each cursor measure, a repeated measures one-way ANOVA was used with planned post-test comparisons. The planned comparisons were as follows: each haptic condition against unassisted, gravity wells against haptic cones, gravity wells against haptic funnels, high-friction targets against haptic funnels and high-friction targets against haptic cones. To correct for multiple comparisons the method of Bonferroni was used. Reported results were the ANOVA p value, (F ratio and degrees of freedom [df]), difference in means between groups for each planned comparison with 95% confidence intervals (95% CI) for the difference between the two group means and statistical significance given multiple comparisons.

5.2.2 Missed-click

The results shown in Figure 5.1 are promising in that each technique reduced the mean number of missed-clicks when compared to the unassisted experiment. A one-way ANOVA shows that the haptic condition had a statistically significant effect on the number of missed-clicks recorded, ($F_{4,20} = 3.7$, $p = 0.02$). However, the results in Table 5.1 show that haptic cones were the only technique to produce a statistically significant improvement in comparison to the unassisted interface. Haptic cones produced the lowest mean number of missed-clicks and reduced the frequency by more than half of that recorded in the unassisted experiment. The confidence intervals for unassisted vs. high-friction targets (95% CI, -0.1 to 4.9) imply that the haptic condition was close to showing a statistically significant improvement.

5.2.3 Missed-click on click

A one way ANOVA suggests that the haptic conditions do not have a significant effect on the mean number of missed-clicks on click ($F_{4,20} = 0.93$, $p = 0.4671$). Only

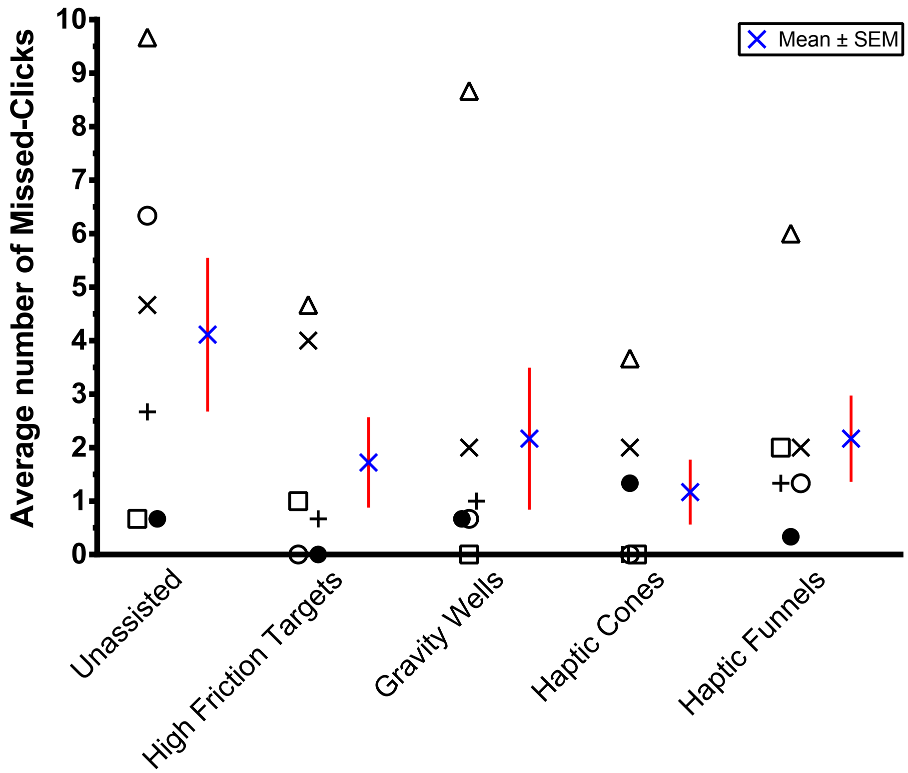

Figure 5.1: The number of missed-clicks recorded over fifty successful selections for each haptic condition with three repetitions.

three participants recorded missed-clicks on click during the unassisted experiment. The low frequency is likely to be because participants tend to slip off targets when clicking rather than slipping onto them. However, the results in Figure 5.2 show that only one participant produced missed-clicks on click when using haptic cones. It is difficult to make definite conclusions using this measure because if the operator does not enter a target region before performing a click then the assistance will not have been given an opportunity to help them.

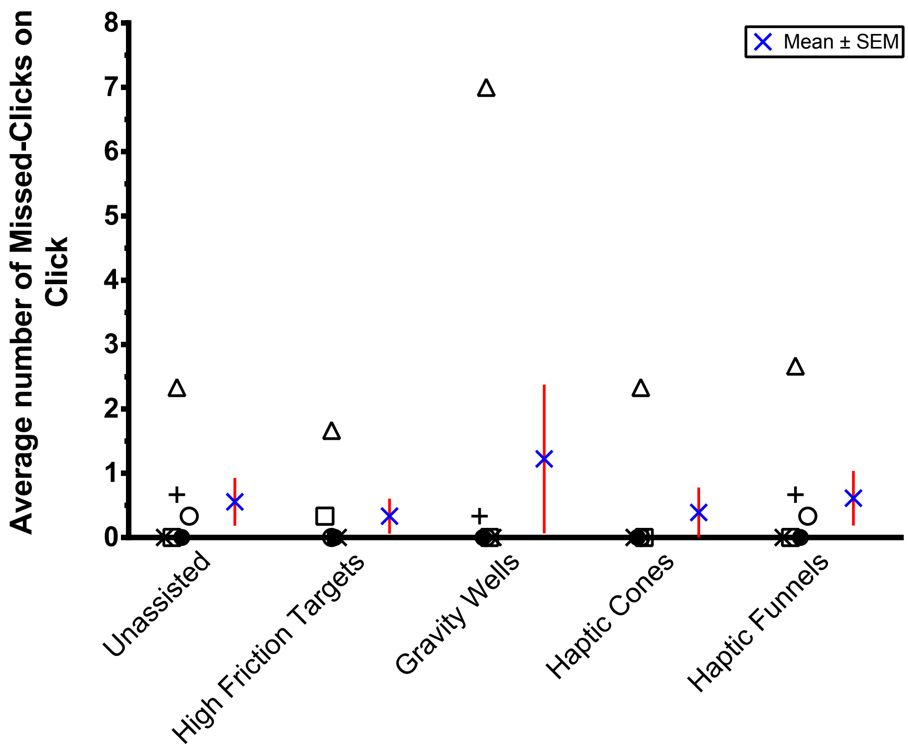

Figure 5.2: The number of missed-clicks on click recorded over fifty successful selections for each haptic condition with three repetitions.

5.2.4 Missed-click on release

The missed-click on release measure may be more useful than the missed-click on click because the operator will have definitely entered the target region and been provided with assistance. The results are presented in Figure 5.3. A one-way ANOVA shows that the haptic condition has a significant effect on the mean number of missed-clicks on release ($F_{4,20} = 3.03$, $p = 0.0419$). None of the planned comparisons reached formal statistical significance in Table 5.1. However, the confidence intervals for unassisted vs. high-friction targets (95% CI, -0.1 to 2.7), unassisted vs. gravity wells (95% CI, -0.2 to 2.6) and unassisted vs. haptic cones (95% CI, 0.0 to 2.8) imply that these techniques were close to showing a statistically significant improvement. Haptic cones reduced the mean number of missed-clicks on release to zero for all participants. The improvement in this measure will be credited to the stability provided to the operator through the clamping at the cone apex.

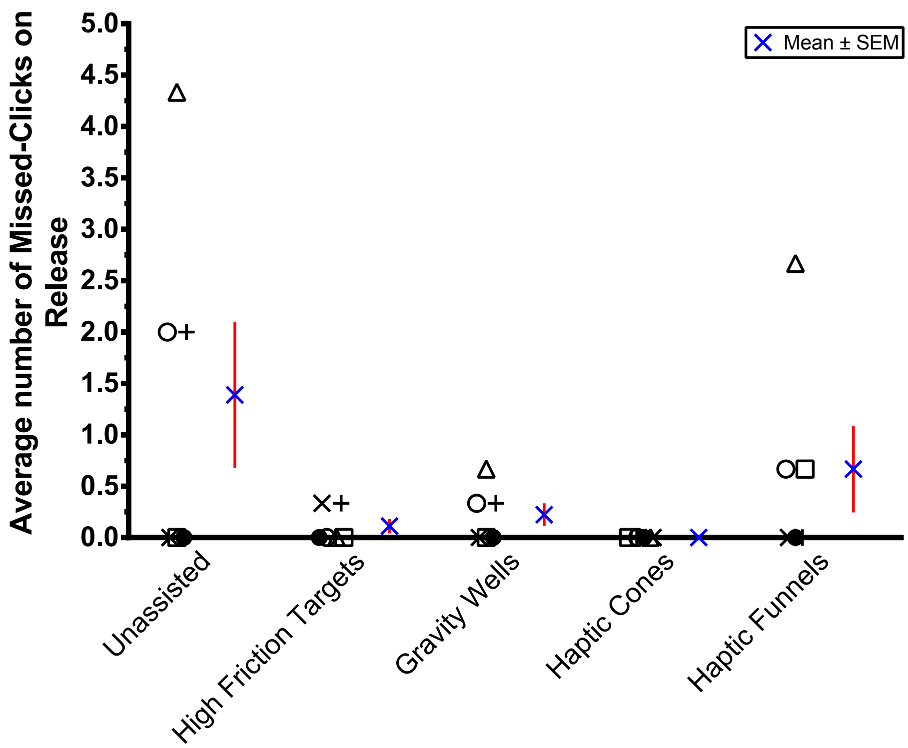

Figure 5.3: The number of missed-clicks on release recorded over fifty successful selections for each haptic condition with three repetitions.

5.2.5 Click-release distance travelled

A one-way ANOVA shows that the haptic condition has a significant effect on the mean click-release distance travelled ($F_{4,20} = 9.49$, $p = 0.0002$). The results in Figure 5.4 show that high-friction targets and haptic cones were the most effective at reducing the click-release distance travelled. Both techniques produced a statistically significant improvement compared to the unassisted experiment, as shown in Table 5.1. The results also show that haptic cones produced a significant improvement compared to gravity wells. The spring force of gravity wells tends to cause the cursor to oscillate slightly about the target centre, which may explain why no significant improvements were recorded. The haptic funnels seem to have little effect on the click-release distance travelled. If the operator does not make an effort to navigate to the walls at the target centre then the assistance may be less effective.

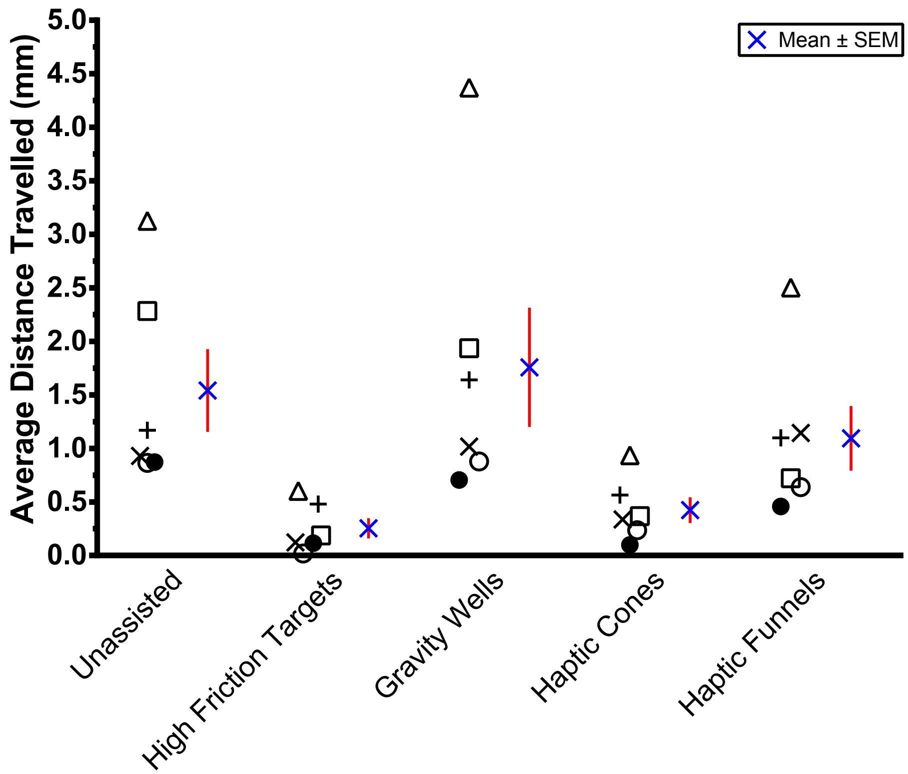

Figure 5.4: The successful click-release distance travelled recorded over fifty successful selections for each haptic condition with three repetitions.

5.2.6 Click-release displacement

A one-way ANOVA shows that the haptic condition has a significant effect on the mean click-release displacement ($F_{4,20} = 10.3$, $p = 0.0001$). The results in Figure

5.5 and the planned comparisons in Table 5.1 show that haptic cones and high-friction targets were the only techniques to produce significant improvements over the unassisted interface. Haptic cones also produced significant improvements in the click-release displacement compared to gravity wells. It was expected that gravity wells would improve clicking stability due to the clamping at the target centre. However, it appears that the spring force may throw the cursor slightly, which is why the distance between the click and release is not as close as expected.

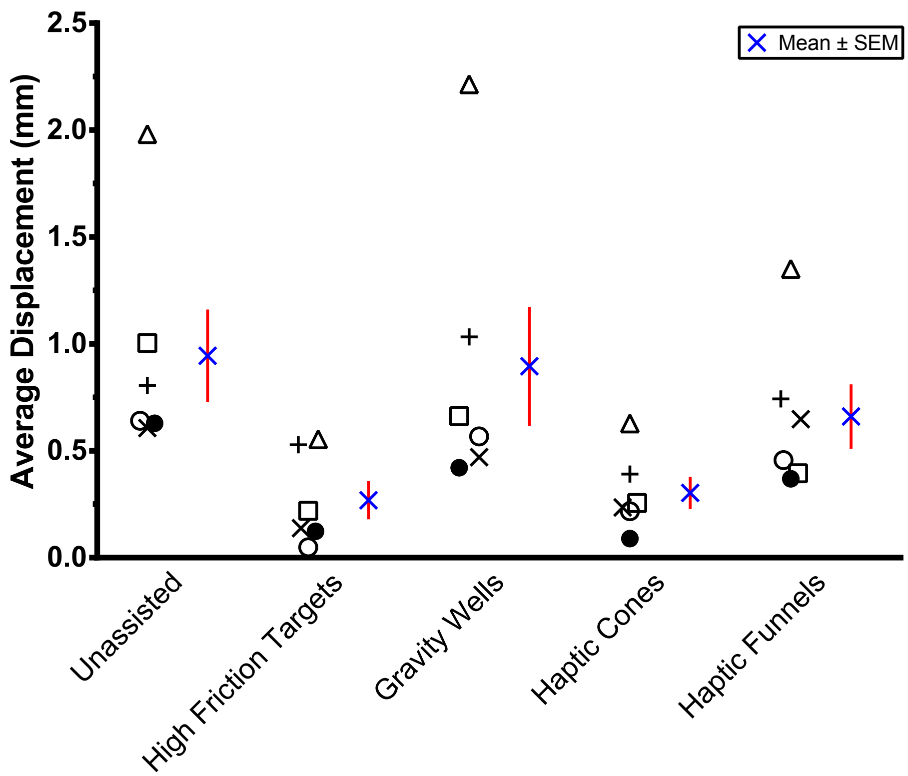

Figure 5.5: The successful click-release displacement recorded over fifty successful selections for each haptic condition with three repetitions.

5.2.7 On-click distance to target centre line

A one-way ANOVA shows that the haptic condition has a significant effect on the mean on-click distance to target centre line ($F_{4,20} = 10.6$, $p < 0.0001$). Haptic cones were the most successful technique at guiding the operator to the target centre, as shown in Figure 5.6. The planned comparisons in Table 5.1 show that haptic cones produced significant improvements over the unassisted interface, gravity wells and high-friction targets. High-friction targets were the least effective technique at reducing the on-click distance to target centre line. It was observed that many of the participants had difficulty drawing the cursor to the target centre due to the friction force experienced at the edges. As a result, participants would often click near the target's edge, which means that any unwanted movements may draw the cursor off the target and result in a missed-click.

5.2.8 Throughput

The haptic assistive techniques in this study are designed to improve clicking accuracy without adversely affecting the speed component. The results for throughput are shown in Figure 5.7. A one-way ANOVA shows that the haptic condition has a significant effect on the mean throughput ($F_{4,20} = 10.5$, $p < 0.0001$). The planned comparisons in Table 5.1 show that haptic cones produced significant improvements over the unassisted interface, gravity wells and high-friction targets. This will be credited to the technique improving clicking accuracy without adversely affecting speed. The results for the three other haptic conditions are not significantly different to the unassisted interface.

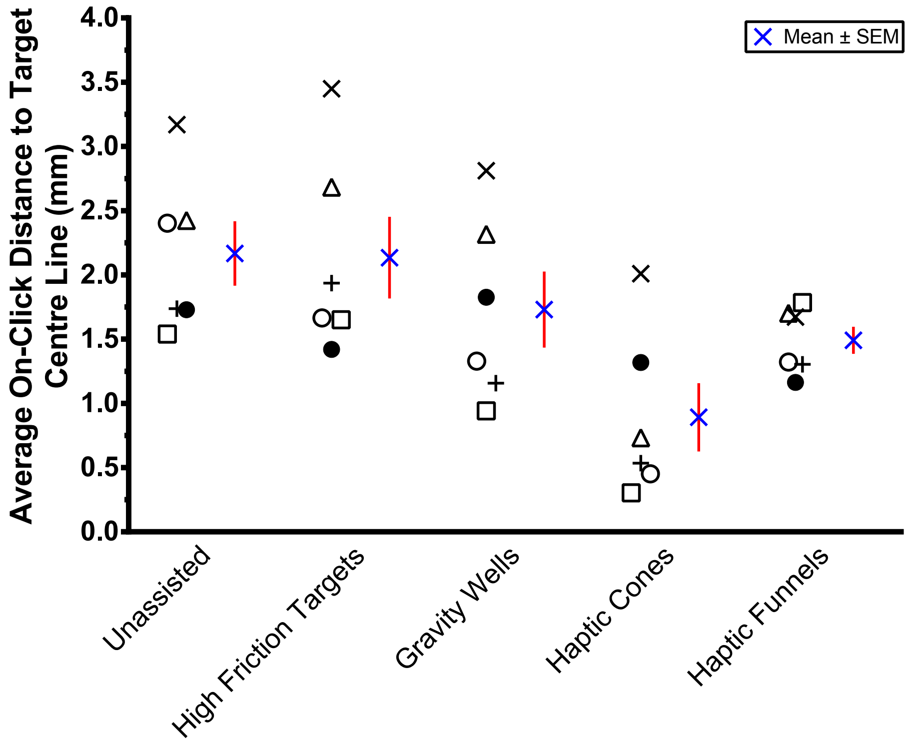

Figure 5.6: The on-click distance to target centre line recorded over fifty successful selections for each haptic condition with three repetitions.

5.2.9 Experiment time

A one-way ANOVA shows that the haptic condition had a statistically significant effect on the mean experiment time, ($F_{4,20} = 8.46$, $p = 0.0004$). The results are presented in Figure 5.8. The planned comparisons in Table 5.1 show that gravity wells were significantly slower than haptic cones and the unassisted interface. This is likely to be caused by the operator having to oppose forces from gravity well target distracters and the subsequent corrections required when exiting them. Haptic funnels were significantly slower than high-friction targets. This suggests that the time penalty associated with exiting funnels in the opposite direction may be greater than scrolling over the surface of high-friction targets.

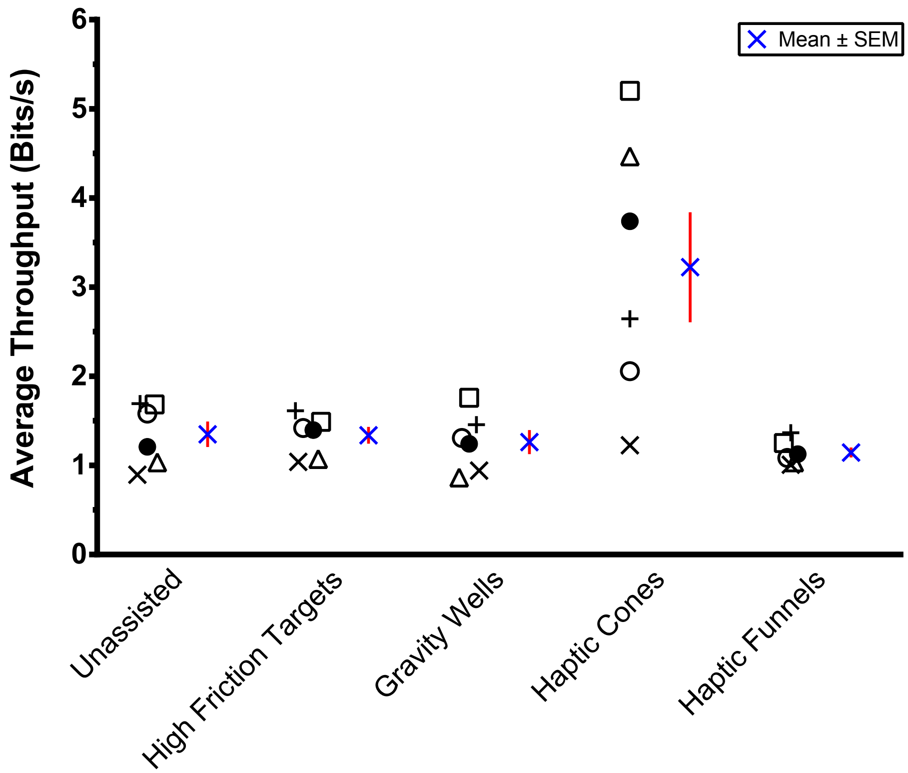

Figure 5.7: The throughput recorded over fifty successful selections for each haptic condition with three repetitions.

5.2.10 Results of multiple comparisons

To provide information on the statistical significance of the improvements a repeated measures one-way ANOVA was performed with planned post-test comparisons. Since the null hypothesis was rejected a Bonferroni multiple comparison was used to determine which means are different. The results are shown in Table 5.1.

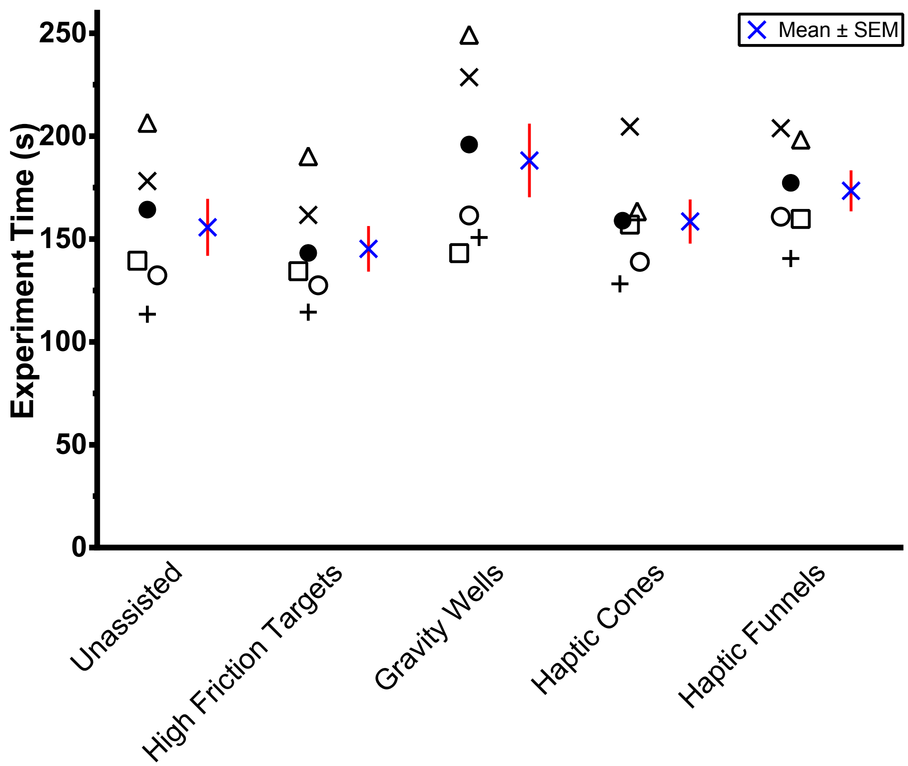

Figure 5.8: The experiment time recorded over fifty successful selections for each haptic condition with three repetitions.

Planned Comparisons			Missed-Clicks	Missed-Click on Click	Missed-Click on Release	Click-Release Distance Travelled (mm)	Click Displacement (mm)	Centre Line (mm)	Throughput (Bits/s)	Experiment time (s)
Unassisted	*vs.*	High Friction Targets	2.4 (-0.1 to 4.9)	0.2 (-1.4 to 1.8)	1.3 (-0.1 to 2.7)	1.3 ** (0.4 to 2.2)	0.7 *** (0.2 to 1.1)	0.0 (-0.7 to 0.7)	0.01 (-1.16 to 1.18)	10.4 (-14.5 to 35.4)
Unassisted	*vs.*	Gravity Wells	1.9 (-0.5 to 4.4)	-0.7 (-2.3 to 0.9)	1.2 (-0.2 to 2.6)	-0.2 (-1.2 to 0.7)	0.0 (-0.4 to 0.5)	0.4 (-0.3 to 1.1)	0.09 (-1.08 to 1.25)	-32.5 ** (-57.4 to -7.58)
Unassisted	*vs.*	Haptic Cones	2.9 * (0.5 to 5.4)	0.2 (-1.4 to 1.8)	1.4 (0.0 to 2.8)	1.1 * (0.2 to 2.1)	0.6 ** (0.2 to 1.1)	1.3 *** (0.6 to 2.0)	-1.87 *** (-3.04 to -0.71)	-2.82 (-27.7 to 22.1)
Unassisted	*vs.*	Haptic Funnels	1.9 (-0.5 to 4.4)	-0.1 (-1.6 to 1.5)	0.7 (-0.7 to 2.1)	0.4 (-0.5 to 1.4)	0.3 (-0.1 to 0.7)	0.7 (0.0 to 1.37)	0.20 (-0.96 to 1.37)	-17.7 (-42.6 to 7.18)
Gravity Wells	*vs.*	Haptic Cones	1 (-1.5 to 3.5)	0.8 (-0.8 to 2.4)	0.2 (-1.2 to 1.6)	1.3 ** (0.4 to 2.3)	0.6 ** (0.2 to 1.0)	0.8 * (0.1 to 1.5)	-1.96 *** (-3.13 to -0.79)	29.7 * (4.76 to 54.6)
Gravity Wells	*vs.*	Haptic Funnels	0.0 (-2.5 to 2.5)	0.6 (-1.0 to 2.2)	-0.4 (-1.9 to 1.0)	0.7 (-0.3 to 1.6)	0.2 (-0.2 to 0.7)	0.2 (-0.5 to 0.9)	0.12 (-1.05 to 1.28)	14.8 (-10.1 to 39.7)
High Friction Targets	*vs.*	Haptic Funnels	-0.4 (-2.9 to 2.0)	-0.3 (-1.9 to 1.3)	-0.6 (-2.0 to 0.9)	-0.8 (-1.8 to 0.1)	-0.4 (-0.8 to 0.0)	0.6 (-0.1 to 1.3)	0.19 (-0.97 to 1.36)	-28.2 * (-53.1 to -3.27)
High Friction Targets	*vs.*	Haptic Cones	0.6 (-1.9 to 3.0)	-0.1 (-1.6 to 1.5)	0.1 (-1.3 to 1.5)	-0.2 (-1.1 to 0.8)	0.0 (-0.5 to 0.4)	1.2 *** (0.5 to 1.9)	-1.88 *** (-3.05 to -0.72)	-13.3 (-38.2 to 11.6)

Table 5.1: Bonferroni post-test multiple comparisons. The reported measures are mean difference, significance levels and (95% confidence intervals). Significance levels are reported as * for ($0.01 < p \leq 0.05$), ** for ($0.001 < p \leq 0.01$), *** for ($0.0001 < p \leq 0.001$) and **** for ($p \leq 0.0001$).

5.2.11 Discussion

Experiment 1 was designed to investigate the performance benefits of haptic assistance that does not require force calibration compared to traditional techniques in a static configuration. The goal was to design haptic assistance that outperforms traditional techniques and alleviates or at least reduces the shortcomings highlighted in Chapter 2. The results show that haptic cones outperformed gravity wells and high-friction targets in terms of reducing error rates during a point-and-click task. Haptic cones reduced the mean number of missed-clicks by 71% and were the only technique to show a statistically significant improvement over the unassisted interface. Many of the cursor measures help to explain why this is the case. For example, the click-release distance travelled and displacement show that the clamping at the apex helps maintain cursor stability when clicking. Haptic cones were the most effective technique at guiding the cursor to the target centre, which is beneficial because if the operator slips it is more likely that the cursor will remain within the target region when the switch is released. No missed-clicks on release were recorded for haptic cones, which will be credited to the effective clamping at the target centre.

The results for V-shaped funnels were not as promising as haptic cones but the performance levels were not dissimilar to the traditional haptic techniques. The lack of significant performance increase is likely to be because haptic funnels are less effective at guiding the cursor to the target centre. If the operator does not make a conscience effort to navigate to the centre of the target then they will receive little assistance from the funnel walls. This would explain why the performance levels are comparatively lower for the click-release distance travelled, click-release displacement and on-click distance to target centre line.

The haptic cones and V-shaped funnels have many other benefits over traditional haptic assistance. One of the major advantages is that they do not require force calibration to optimise interaction. As a result, the operator can use the assistance immediately and any changes in their clicking characteristics, such as fatigue, will not require force recalibration. Even if traditional haptic techniques could be calibrated to suit a person's cursor movement characteristics there are no guarantees that the assistance will be beneficial in the long term. For example, if an individual's impairment deteriorates then the system needs to detect this and react accordingly, both of which are not trivial tasks due to the uniqueness of disability. The unpredictable data inputted by people with physical impairments can make it difficult to effectively tune methods of interaction.

The second advantage of haptic cones and V-shaped funnels is that the operator does not have to oppose a force to exit target distracters. Traditional techniques, such as gravity wells and high-friction targets, impose a spring or frictional force on the operator that they have to overcome before exiting. Haptic cones can be exited by simply navigating the cone walls and haptic funnels can be exited in the opposite direction to which they were first entered. The ability to exit haptic cues easily means that target distracters are less intrusive on interaction. People with decreased muscle strength or joint difficulties will be able to use the assistance without discomfort from imposing forces, which means they are more suitable for a wider range of impairments. Previously, the "snap effect" of gravity wells has been a concern for people with joint difficulties or decreased muscle strength.

A number of studies have highlighted that the steepest learning curve tends to be early on in the use of a pointing device [DM94] [DKM99] [MKS01]. It is generally accepted that a block of 10 trials similar to those discussed in Section 4.2.3 is sufficient enough to negate any learning effects. Typically, these studies use the experiment time as a benchmark to determine if there are any learning effects or issues regarding fatigue. To ensure that there was not an ordering effect between the 3 repetitions a repeated measures one-way ANOVA was performed to test the null hypothesis that there were performance differences between the repetitions for the experiment time and number of missed-clicks for each haptic condition. The results show that no statistically significant differences were recorded for the experiment time between repetitions for **unassisted** ($F_{2,10} = 0.08$, $p = 0.926$), **gravity wells** ($F_{2,10} = 0.77$, $p = 0.4894$), **high-friction targets** ($F_{2,10} = 0.02$, $p = 0.9838$), **haptic cones** ($F_{2,10} = 0.47$, $p = 0.6405$) or **haptic funnels** ($F_{2,10} = 3.36$, $p = 0.0765$). No statistically significant differences were recorded for the number of missed-clicks between repetitions for **unassisted** ($F_{2,10} = 2.58$, $p = 0.1246$), **gravity wells** ($F_{2,10} = 1.87$, $p = 0.2042$), **high-friction targets** ($F_{2,10} = 1.46$, $p = 0.2774$), **haptic cones** ($F_{2,10}$

= 0.97, $p = 0.4124$) or **haptic funnels** ($F_{2,10} = 0.47$, $p = 0.6401$). Taken together, these results indicate that there are no significant learning effects in terms of speed or accuracy for each haptic condition and confirm that the duration between repetitions is significant enough to ensure that fatigue is not a factor.

5.2.12 Conclusion

This section has presented a study of pointing device accuracy for six motion-impaired participants using four separate haptic conditions. The Windows OSK was chosen as a realistic interface to investigate the performance benefits of haptic assistance during a point-and-click task. Haptic cones were the most effective technique for decreasing the number of missed-clicks and improving throughput, the two measures of performance that are most important. The new cursor measures presented in this thesis have been useful in identifying why clicking errors occur and how haptic techniques may assist in these areas. Many of the shortcomings highlighted in Chapter 2 have been alleviated by utilising a 3DOF interface. By overlaying the partially transparent window on top of the main OpenGL viewport it has been possible to provide suitable visual cues to accompany the haptic conditions.

5.3 Experiment 2

5.3.1 Experimental procedure

Experiment 2 investigates the performance benefits of virtual switches, deformable switches and deformable cones. The aim of these techniques is to further improve the error rates and targeting times reported in Experiment 1. Relevant data for gravity wells and haptic cones has been included from the previous experiment. Gravity wells have been included in this study because they are the most widely reported of all haptic assistance. Haptic cones have been included because they were the most effective technique for reducing the number of missed-clicks and improving throughput. These will both be useful in providing a comparison to the deformable techniques and virtual switches. Experiment 2 also concentrates on the intrusiveness of haptic assistance and the effect that target distracters have on interaction. The previous experiment highlighted the potential that haptic cones have for improving interaction rates but did not provide a detailed insight into the effect of distracters. One of the limitations of haptic cones is that they are in operation at all times, for all icons, which can make it difficult to smoothly scroll across the screen whilst remaining in contact with the virtual plane, as demonstrated in Video 6 http://youtu.be/cytiXh3aXFs. Although distracters are easier to exit, it would be desirable for the operator to have the choice of entering them or not. Another limitation of the current haptic cone approach is that the operator is not given any assistance to exit a target. Therefore, they have to either manually navigate the cone wall or pull out of the cone before performing the next operation. The deformable techniques proposed in this thesis aim to extend the haptic cone approach to overcome these shortcomings.

The six participants were asked to perform the on-screen keyboard (OSK) task described in Section 4.2.2. Each operator was required to create a predefined sentence within a textbox using the OSK. The active target was always highlighted in red and only progressed to the next key once acquired. A sequence comprised of 50 successful target selections and data collection began after the first target had been selected. Data collection was continuous within a sequence and breaks of a minimum of 10 minutes were taken between trials. In accordance with Fitts' law participants were instructed to work as quickly as possible whilst maintaining a high level of accuracy. Due to only having a small sample size the experiments were repeated three times for each haptic condition. The order of presentation of the haptic techniques was randomised for each person. For the unassisted condition the force feedback provided by the haptic cues was turned off, meaning the Phantom Omni operated as an "ordinary" pointing device. Sessions were two hours long and shared between the user group. The cursor movements of each participant were recorded using the data recorder discussed in Section 4.2.5 and analysed according to the measures described previously in Sections 2.7 and 3.2.

Fitts' law is often used as the model for cursor movement in HCI. It can be used to give a measure of the trade-off between speed and accuracy, known as throughput (**TP**). However, the haptic conditions investigated in this section are very effective at guiding the cursor to the target centre, which makes it difficult to compute throughput. Although the overall average on-click distance to target centre line may not be zero, there were many occasions where all the participants clicked at the target centre for all repetitions. As a result, there is often no standard deviation in the click positions and it is not possible to calculate the index of difficulty required for throughput. Therefore, the main emphasis of the cursor analysis in this section is on error rates, targeting times and the effects of distracters.

The data has been described for each haptic condition using the mean of the 3 repetitions for each participant and standard error of the mean (SEM). The participants' physical background and corresponding legend can be found in Table 4.2. For each cursor measure (apart from experiment time spent on the virtual plane) a reduction in magnitude is desirable and will signify an improvement. To test the null hypothesis that there were no differences between haptic conditions for each cursor measure, a repeated measures one-way ANOVA was used with planned post-test comparisons. The planned comparisons were as follows: each haptic condition against unassisted, gravity wells against deformable cones, haptic cones against deformable cones, gravity wells against deformable switches and virtual switches against deformable switches. To correct for multiple comparisons the method of Bonferroni was used. Reported results were the ANOVA p value, (F ratio and degrees of freedom [df]), difference in means between groups for each planned comparison with 95% confidence intervals (95% CI) for the difference between the two group means and statistical significance given multiple comparisons.

5.3.2 Missed-click

A one-way ANOVA shows that the haptic condition had a statistically significant effect on the number of missed-clicks recorded, ($F_{5,25} = 4.78$, $p = 0.0033$). The results shown in Figure 5.9 are promising in that deformable cones, virtual switches and deformable switches all produced a significant reduction in the mean number of missed-clicks when compared to the unassisted interface. The statistical significance is confirmed in Table 5.2. Deformable switches produced the lowest mean number of missed-clicks closely followed by virtual switches and deformable cones. The deformable techniques were shown to reduce the frequency by more than 75% of that recorded in the unassisted experiment.

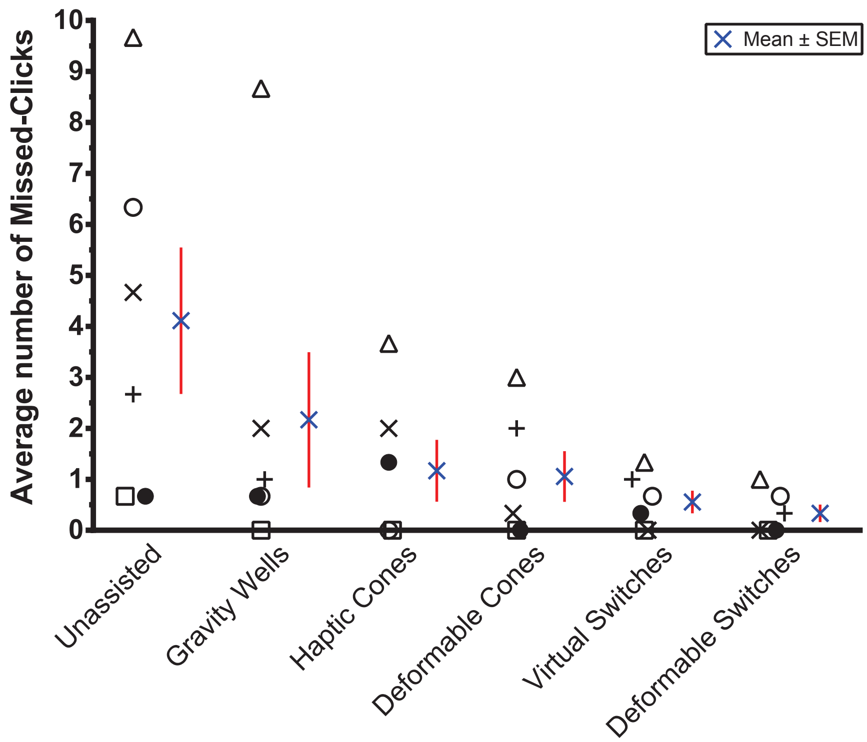

Figure 5.9: The number of missed-clicks recorded over fifty successful selections for each haptic condition with three repetitions.

5.3.3 Click-release displacement

A one-way ANOVA shows that the haptic condition had a statistically significant effect on the mean click-release displacement, ($F_{5,25} = 7.63$, $p = 0.0002$). Figure 5.10 and Table 5.2 show that haptic cones, deformable cones and virtual switches all significantly improved the click-release displacement compared to the unassisted interface. Deformable switches reduced the mean click-release displacement compared to the unassisted experiment but did not reach formal statistical significance in the planned comparisons. However, the confidence intervals for unassisted vs. deformable switches (95% CI, -0.1 to 0.8) suggest that the haptic condition was close to showing a statistically significant improvement. This is less of an issue for the virtual switches because the release will be recorded when the spring is restored regardless of whether the cursor is accurately positioned inside the target or not. The click-release displacement measure is most useful for comparing techniques that require the device switch to be operated.

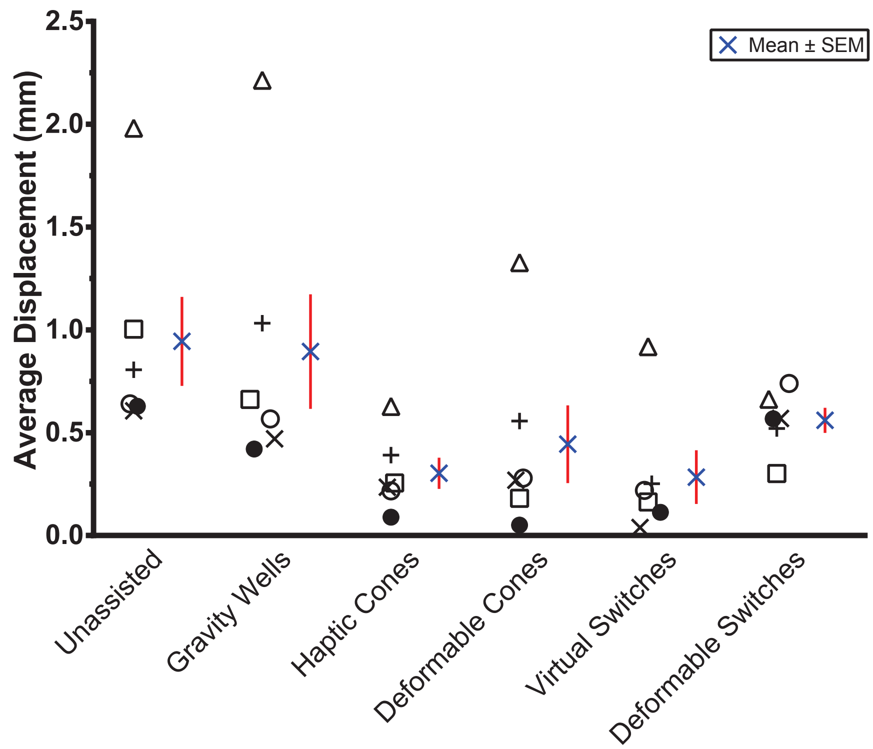

Figure 5.10: The successful click-release displacement recorded over fifty successful selections for each haptic condition with three repetitions.

5.3.4 On-click distance to target centre line

A one-way ANOVA shows that the mean on-click distance to target centre line differed significantly across the haptic conditions, ($F_{5,25} = 16.1$, $p < 0.0001$). Gravity wells were the only haptic condition to not show significant improvements in guiding the cursor to the target centre line when compared to the unassisted experiment, as shown in Figure 5.11 and Table 5.2. Deformable switches produced the most significant improvement, closely followed by virtual switches and deformable cones. Both deformable techniques were significantly more effective at guiding the operator to the centre of the target than traditional gravity wells. Accurate selection at the target centre is important because if the operator slips it is more likely that the cursor will remain within the target region when the switch is released. Unwanted movements near the edges may draw the cursor off the target and result in a missed-click.

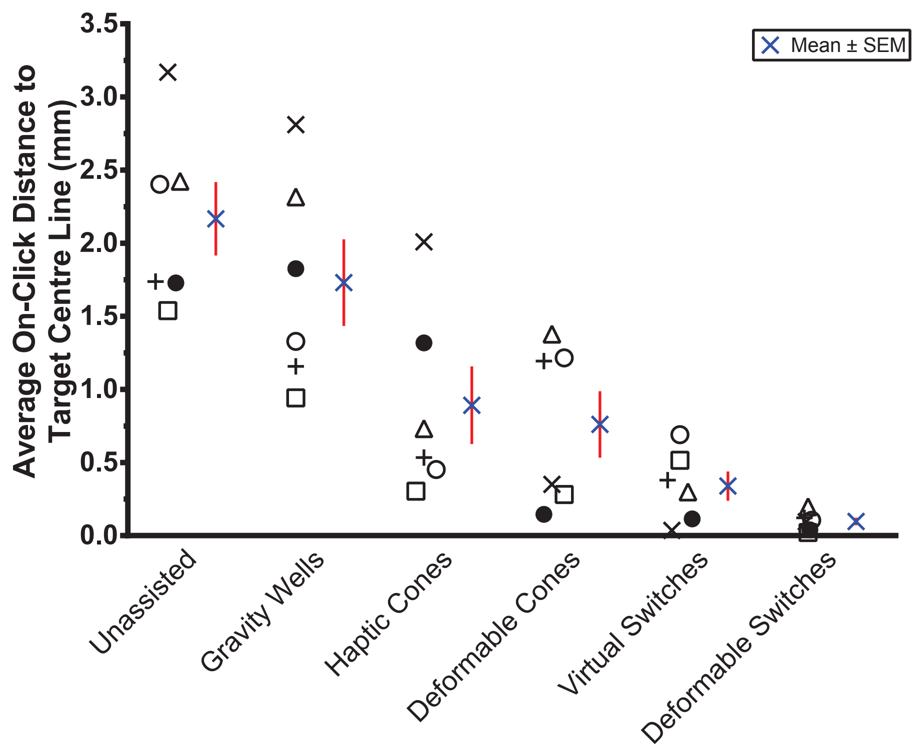

Figure 5.11: The on-click distance to target centre line recorded over fifty successful selections for each haptic condition with three repetitions.

5.3.5 Percentage of experiment time spent on the virtual plane

The amount of time spent in contact with the virtual plane will give an indication of how intrusive each haptic technique is on interaction. The unassisted experiment will give a benchmark for the natural amount of time that participants spend in contact with the virtual plane. A one-way ANOVA showed that the effect of the haptic condition was significant on the amount of time spent in contact with the virtual plane, ($F_{5,25} = 22.3$, $p < 0.0001$). Table 5.2 shows that the participants spent significantly less time in contact with the virtual plane for gravity wells, haptic cones and virtual switches in comparison to the unassisted experiment. This will be a result of participants lifting the proxy off the virtual plane because they were not able to easily pass over distracters. It is clear from the results in Figure 5.12 and Table 5.2 that a much greater percentage of experiment time was spent in contact with the virtual plane for the two deformable techniques in comparison to the other assistance. No statistically significant differences were recorded for the deformable techniques compared against the unassisted interface. Taken together, these results suggest that the deformable techniques are much less intrusive on interaction because the operator decided to pull off the virtual plane less often to pass over potential distracters. This will be credited to the fact that the operator can choose when to use or ignore the assistance.

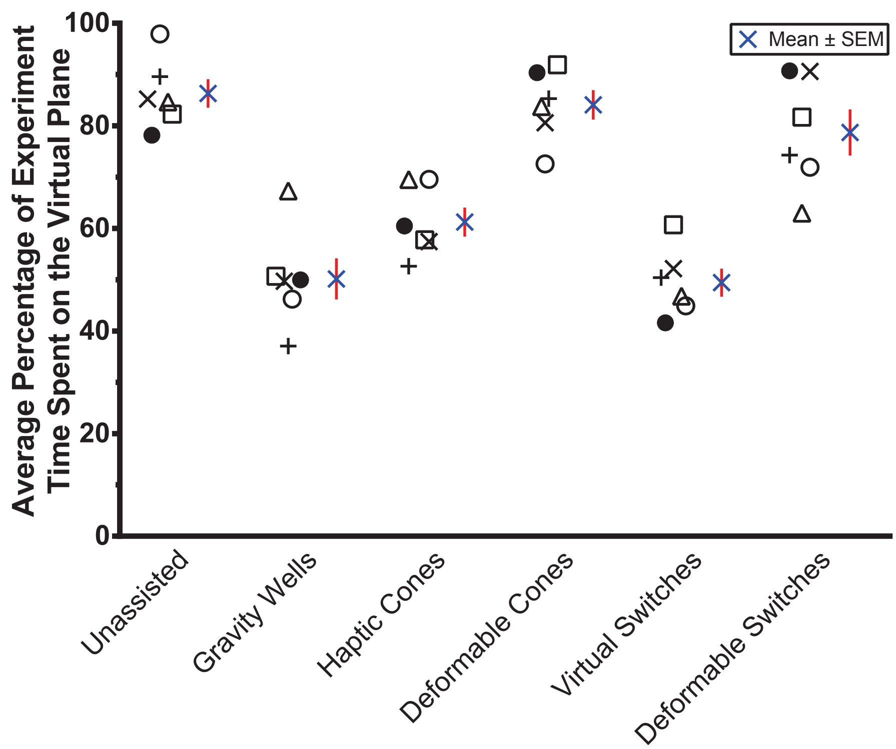

Figure 5.12: The percentage of experiment time spent on the virtual plane recorded over fifty successful selections for each haptic condition with three repetitions.

5.3.6 Experiment distance travelled

Ideally the operator will move the shortest distance over the course of an experiment, i.e. they will take the most direct routes. If the experiment distance travelled for a given haptic condition is significantly worse than the unassisted experiment then this will indicate that the operator had to make further corrections due to the intrusion of target distracters. A one-way ANOVA shows that the difference in experiment distance travelled between the haptic conditions was statistically significant, ($F_{5,25}$ $= 28.4$, $p < 0.0001$). It was observed that participants often travelled further when gravity wells were in operation due to the corrections required when exiting and overshooting distracters. The results presented in Figure 5.13 suggest that intrusive techniques that are in operation at all times result in the cursor travelling a greater distance over the course of the experiment. This is confirmed in Table 5.2 where the distance travelled for gravity wells and virtual switches was significantly greater than the unassisted interface. The deformable techniques allow the assistance to be ignored more easily and do not show significant differences in the distance travelled when compared to the unassisted experiment.

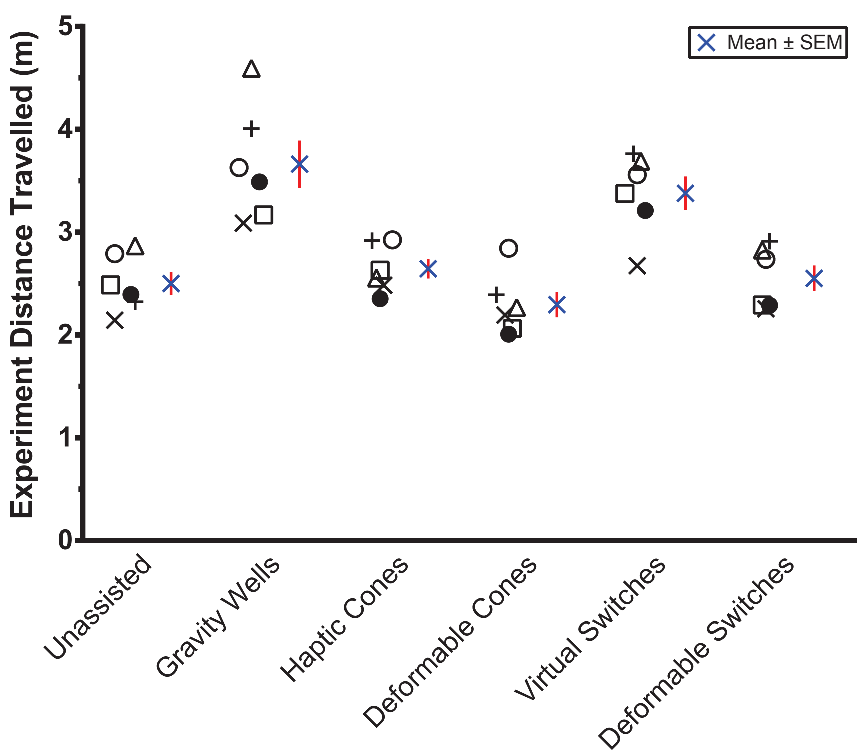

Figure 5.13: The experiment distance travelled recorded over fifty successful selections for each haptic condition with three repetitions.

5.3.7 Experiment time

It has been observed that the experiment time often increases when intrusive haptic techniques are used in conjunction with densely populated interfaces. This is mainly due to the operator having to oppose forces from target distracters. A one-way ANOVA shows that the haptic condition had a statistically significant effect on the mean experiment time, ($F_{5,25} = 11.5$, $p < 0.0001$). The results in Figure 5.14 show that the most significant reductions in experiment time occur when using assistance that can be easily ignored. This is confirmed in Table 5.2 where deformable cones and deformable switches were the only techniques to show statistically significant improvements compared to the unassisted experiment. This will be credited to the ease of navigating to the target and the increased confidence in selecting it. All the other techniques that are in operation at all times either worsened the experiment time or had no significant effect. Both deformable techniques produced significant improvements in experiment time compared to the traditional gravity wells.

5.3.8 Results of multiple comparisons

To provide information on the statistical significance of the improvements a repeated measures one-way ANOVA was performed with planned post-test comparisons. Since the null hypothesis was rejected a Bonferroni multiple comparison was used to determine which means are different. The results are shown in Table 5.2.

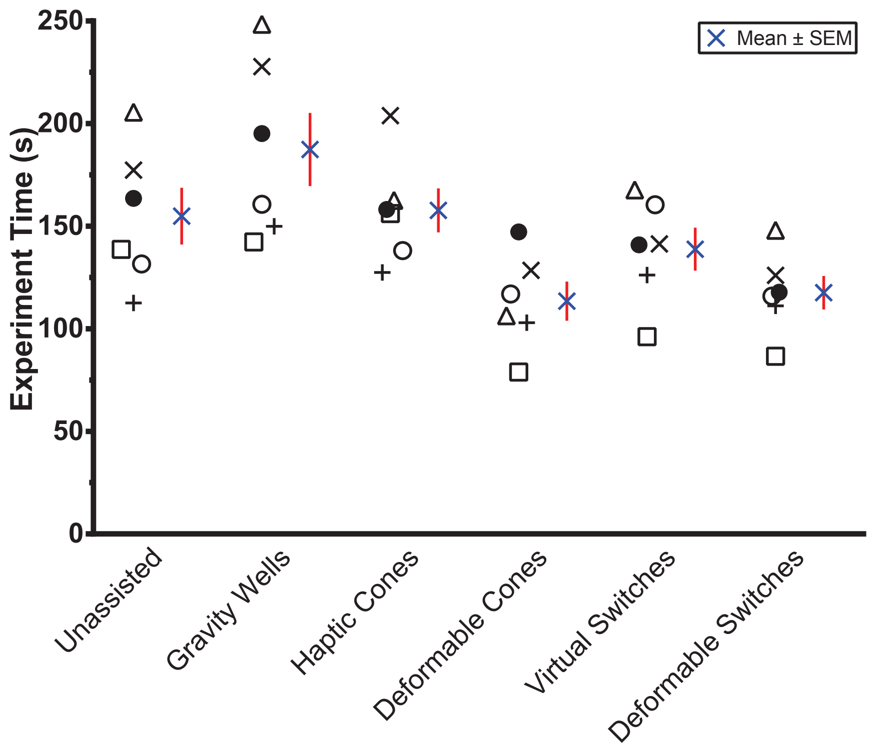

Figure 5.14: The experiment time recorded over fifty successful selections for each haptic condition with three repetitions.

Planned Comparisons			Missed-Clicks	Click Displacement (mm)	Centre Line (mm)	% Time on Virtual Plane	Distance Travelled (m)	Experiment Time (s)
Unassisted	*vs.*	Gravity Wells	1.9 (-0.8 to 4.7)	0.0 (-0.4 to 0.5)	0.4 (-0.4 to 1.3)	36.1 **** (20.9 to 51.4)	-1.16 **** (-1.60 to -0.72)	-32.5 (-67.6 to 2.6)
Unassisted	*vs.*	Haptic Cones	2.9 * (0.2 to 5.7)	0.6 ** (0.2 to 1.1)	1.3 ** (0.4 to 2.1)	25.1 *** (9.8 to 40.4)	-0.14 (-0.58 to 0.30)	-2.9 (-38.0 to 32.3)
Unassisted	*vs.*	Deform. Cones	3.1 * (0.3 to 5.8)	0.5 * (0.1 to 0.9)	1.4 *** (0.5 to 2.3)	2.2 (-13.1 to 17.5)	0.206 (-0.23 to 0.65)	41.4 * (6.3 to 76.5)
Unassisted	*vs.*	Switches	3.6 ** (0.8 to 6.3)	0.7 ** (0.2 to 1.1)	1.8 **** (1.0 to 2.7)	36.9 **** (21.6 to 52.2)	-0.88 **** (-1.32 to -0.44)	16.1 (-19.1 to 51.2)
Unassisted	*vs.*	Deform. Switches	3.8 ** (1.0 to 6.5)	0.4 (-0.1 to 0.8)	2.1 **** (1.2 to 2.9)	7.6 (-7.7 to 22.9)	-0.50 (-0.49 to 0.39)	37.3 * (2.1 to 72.4)
Gravity Wells	*vs.*	Deform. Cones	1.1 (-1.6 to 3.9)	0.5 * (0.0 to 0.9)	1.0 * (0.1 to 1.8)	-33.9 **** (-49.2 to -18.6)	1.37 **** (0.93 to 1.81)	73.9 **** (38.8 to 109)
Haptic Cones	*vs.*	Deform. Cones	0.1 (-2.6 to 2.9)	-0.1 (-0.6 to 0.3)	0.1 (-0.7 to 1.0)	-22.9 ** (-38.1 to -7.6)	0.35 (-0.91 to 0.79)	44.3 ** (9.1 to 79.4)
Gravity Wells	*vs.*	Deform. Switches	1.8 (-0.9 to 4.6)	0.3 (-0.1 to 0.8)	1.6 **** (0.8 to 2.5)	-28.5 **** (-43.8 to -13.3)	1.11 **** (0.67 to 1.55)	69.8 **** (34.6 to 105)
Switches	*vs.*	Deform Switches	0.2 (-2.5 to 3.0)	-0.3 (-0.7 to 0.2)	0.2 (-0.6 to 1.1)	-29.3 **** (-44.6 to -14.0)	0.83 **** (0.39 to 1.27)	21.2 (-13.9 to 56.3)

Table 5.2: Bonferroni post-test multiple comparisons. The reported measures are mean difference, significance levels and (95% confidence intervals). Significance levels are reported as * for ($0.01 < p \leq 0.05$), ** for ($0.001 < p \leq 0.01$), *** for ($0.0001 < p \leq 0.001$) and **** for ($p \leq 0.0001$).

5.3.9 Discussion

Experiment 2 was designed to explore the benefits of non-intrusive haptic assistance that the operator can choose to use or ignore. The goal was to design techniques that outperform those investigated in Experiment 1 and to further reduce the effect that target distracters have on interaction. The results of the study show that deformable haptic cones and deformable virtual switches were the most effective techniques for improving targeting times and reducing error rates during a realistic point-and-click task. Deformable cones and deformable switches reduced the mean number of missed-clicks by 75% and 92% respectively in comparison to the unassisted experiment. Many of the cursor measures help to explain why this is the case. For example, the click-release displacement shows that deformable cones are effective at clamping the cursor to the apex, which helps maintain cursor stability when clicking. The on-click distance to target centre line improved significantly, which is beneficial because if the operator slips then it is more likely that the cursor will remain within the target region when the switch is released. Deformable haptic cones and deformable virtual switches both significantly improved the experiment time by 27% and 25% respectively in comparison to the unassisted interface. Both deformable techniques were significantly more effective at reducing the experiment time in comparison to traditional gravity wells. This will be credited to the increased confidence during target selection without the intrusion from neighbouring haptic cues.

The cursor measures confirm that target distracters are less intrusive for deformable cones and deformable switches. The deformable techniques allow the operator to remain in contact with the virtual plane without positional disruption to the cursor from potential distracters. This is confirmed by the fact that there were no significant differences in the percentage of experiment time spent on the virtual plane when compared to the unassisted experiment. The mean percentage of experiment time spent in contact with the virtual plane was 86% for the unassisted interface compared to 84% and 79% for deformable cones and deformable switches respectively. Techniques that are in operation at all times, such as gravity wells, haptic cones and virtual switches, are less easy to ignore and therefore participants often spend significantly less time in contact with the virtual plane. If the operator has to lift off more often to pass over distracters then this will increase the physical workload and disrupt interaction. Both deformable techniques produced significant improvements in the percentage of time spent in contact with the virtual plane compared to their non-deformable counterparts. The experiment distance travelled for deformable cones and deformable switches was similar to that recorded in the unassisted experiment, which suggests that fewer corrections were required for distracters. In contrast, gravity wells and virtual switches significantly worsened the distance the cursor travelled, which is likely to be caused by the corrections required when passing through distracters or by pulling off the virtual plane and re-applying the HIP.

Significant improvements have been recorded in this section despite the fact that the task was performed on a densely populated interface with potential distracters surrounding the target. The levels of statistical significance reported for a small sample size suggest that the haptic conditions have a large effect on user performance.

5.3.10 Conclusion

This section has presented a study of incorporating haptic assistance into existing user interfaces for motion-impaired participants using five separate haptic conditions. The Windows OSK was chosen as a realistic interface to investigate the performance of the haptic conditions during a point-and-click task.

The deformable techniques proposed in this thesis were the most effective at decreasing the number of missed-clicks and improving targeting times, the two measures of performance that are most important. The new cursor measures have been useful in identifying the effects that target distracters have on interaction. Many of the shortcomings highlighted in Section 2.6 have been alleviated by utilising the 3DOF interface to produce assistance that can be easily used or ignored. The deformable techniques allow the operator to navigate more freely without imposing forces from neighbouring haptic cues, as demonstrated in Videos 10 http://youtu.be/GCxWCYLFUN0 and 14 http://youtu.be/qGWvWwBCfuU. The ability to perform a gesture that enables the assistance will allow haptic feedback to be integrated with existing graphical user interfaces without intrusion from target distracters. This is important because it means that distracters are no longer the limiting factor in the development of haptic assistance. It is anticipated that the results produced in this study will be useful in providing assistance that could significantly improve access to existing computer software. By overlaying the partially transparent window on top of the main OpenGL viewport it has been possible to provide suitable visual cues to accompany the haptic conditions.

5.4 Experiment 3

5.4.1 Experimental procedure

The results from the previous sections have shown that deformable cones and deformable switches are the most effective techniques for improving interaction rates and reducing the effects of target distracters. For haptic assistance to be integrated with real-world GUIs it needs to be generalisable for different target sizes and shape. Previous studies have often reported that the greatest improvements in performance have been observed when the assistance is used in conjunction with smaller sized targets [WWBH97] [CB05]. The aim of Experiment 3 is to investigate the effect that target size and shape have on the performance of the two deformable techniques.

The ISO 9241-9 multidirectional point-and-click task, described in Section 4.2.1, has been chosen for this experiment because there is more flexibility to alter the target shape and size than with the OSK. Target distracters are no longer a concern and so the interface will be suitable to evaluate the assistance. Eight targets are uniformly positioned around a circular layout with a diameter of 50mm. The participant is required to first click on the top target, then on the target directly opposite, then the next target clockwise in the sequence and so on around the circular layout. The active target is always highlighted in red and only progresses to the next target once acquired. Data collection was continuous within a sequence and breaks of a minimum of 10 minutes were taken between trials. In accordance with Fitts' law participants were instructed to work as quickly as possible whilst maintaining a high level of accuracy. The experiment was repeated for each type of assistance and each target shape and size. The order of presentation of the technique, target size and target shape was randomised. For the unassisted condition the force feedback provided by the haptic cues was turned off, meaning the Phantom Omni operated as an "ordinary" pointing device. Sessions were two hours long and shared between the user group. The cursor movements of each participant were recorded using the data recorder discussed in Section 4.2.5 and analysed using the experiment time and number of missed-clicks. The size of the targets were categorised as follows: [**small** (3.5mm × 3.5mm), (7mm × 3.5mm), (3.5mm × 7mm)] , [**medium** (7mm × 7mm), (14mm × 7mm), (7mm × 14mm)] , [**large** (14mm × 14mm), (28mm × 14mm), (14mm × 28mm)]. The shape of the targets were categorised as follows: [**square**, **wide rectangles**, **tall rectangles**]. The data in Figure 5.15 has been described for each target size using the mean and standard error of the mean (SEM). The shape is categorised into three separate columns for each haptic condition. Data for each participant is also presented. The participants' physical background and corresponding legend can be found in Table 4.2. Please note that the participant denoted by ● was unavailable for this experimental task.

5.4.2 Target size

To test the null hypothesis that the target size does not have a significant effect on either the number of missed-clicks or the experiment time, a repeated measures one-way ANOVA was used with planned post-test comparisons. The planned comparisons were deformable cones against unassisted and deformable switches against unassisted. To correct for multiple comparisons the method of Bonferroni was used. Reported results were the ANOVA p value, (F ratio and degrees of freedom [df]), difference in means between groups for each planned comparison with 95% confidence intervals (95% CI) for the difference between the two group means and statistical significance given multiple comparisons.

The results show that the haptic condition had a significant effect on the mean number of missed-clicks for **small** ($F_{2,28} = 14.26$, $p < 0.0001$), **medium** ($F_{2,28} = 7.18$, $p = 0.003$) and **large** ($F_{2,28} = 4.59$, $p = 0.0189$) sized targets. The results from Figure 5.15 and Table 5.3 confirm that haptic assistance provides the most significant performance increase when used in conjunction with smaller targets. A statistically significant improvement in the number of missed-clicks was observed for all target sizes for both deformable techniques compared to the unassisted interface. Table 5.3

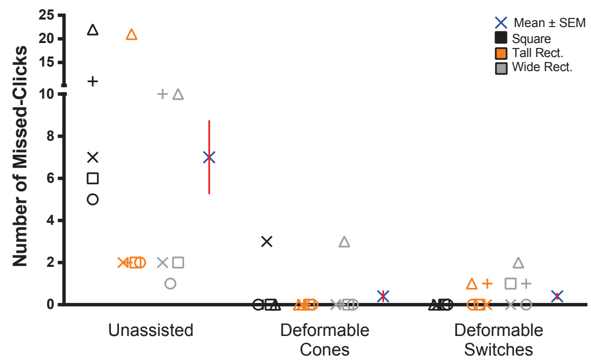

(a) The number of missed-clicks recorded for small sized targets.

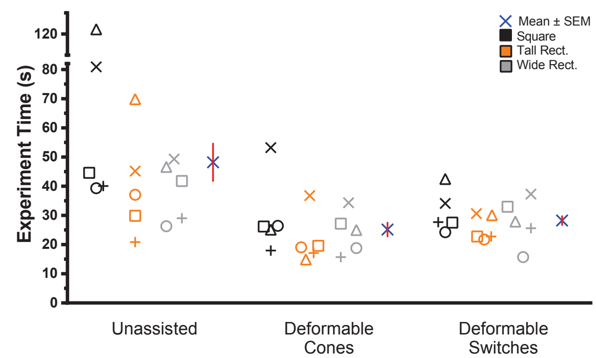

(b) The experiment time recorded for small sized targets.

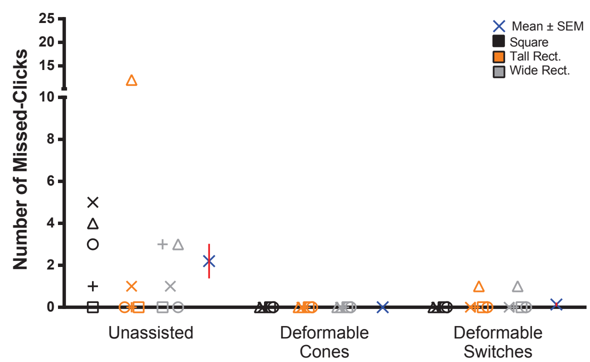

(c) The number of missed-clicks recorded for medium sized targets.

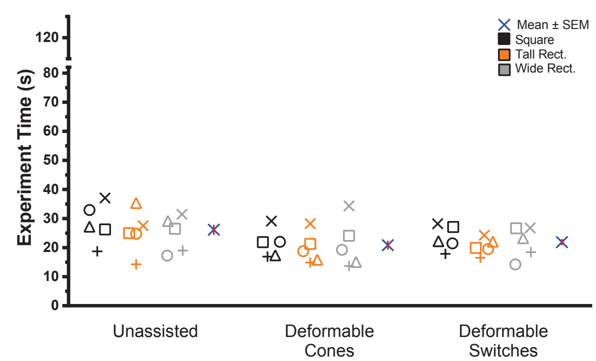

(d) The experiment time recorded for medium sized targets.

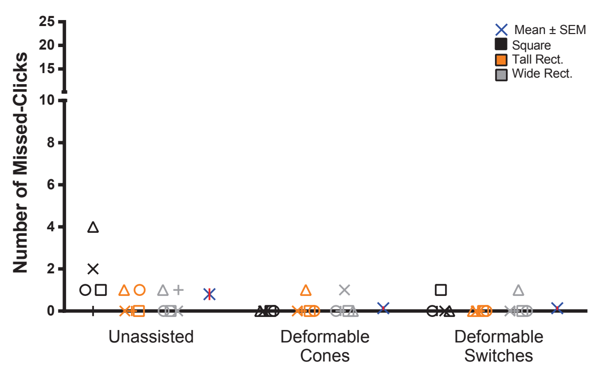

(e) The number of missed-clicks recorded for large sized targets.

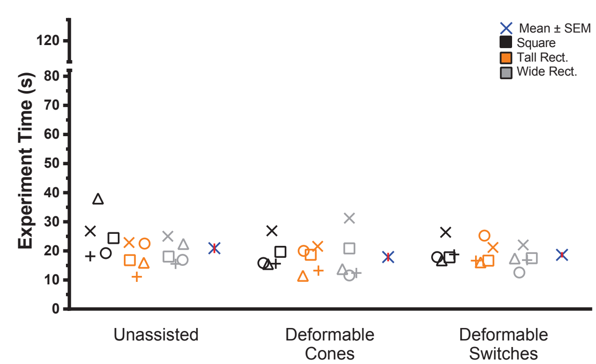

(f) The experiment time recorded for large sized targets.

Figure 5.15: The effect of target size on the performance of each haptic condition. The size of the targets are categorised as small, medium and large.

shows that the level of significance increased as the target size decreased. This is as expected, because smaller targets are more difficult to select and so the assistance will have a more profound effect compared to an unassisted interface.

The results show that the haptic condition had a significant effect on the experiment time for **small** ($F_{2,28} = 13.12$, $p < 0.0001$) and **medium** ($F_{2,28} = 9.06$, $p = 0.0009$) sized targets. Table 5.3 shows that both deformable techniques produced a statistically significant improvement in experiment time for the small and medium sized targets when compared to the unassisted experiment. No significant differences were recorded for **large** ($F_{2,28} = 2.35$, $p = 0.1138$) targets. This is not unexpected, since larger targets tend to be easier to select and so the assistance will have less of an effect on targeting times as the target size increases.

Target Size	Comparison			Missed-Clicks	Experiment Time
Small	Unassisted	*vs.*	Deform. cones	6.6 *** (3.2 to 10.0)	23.0 *** (11.7 to 34.4)
	Unassisted	*vs.*	Deform. switches	6.6 *** (3.2 to 10.0)	20.0 *** (8.6 to 31.3)
Medium	Unassisted	*vs.*	Deform. cones	2.2 ** (0.7 to 3.7)	5.3 *** (2.2 to 8.4)
	Unassisted	*vs.*	Deform. switches	2.1 ** (0.5 to 3.6)	4.3 ** (1.2 to 7.4)
Large	Unassisted	*vs.*	Deform. cones	0.7 * (0.1 to 1.3)	3.0 (-0.4 to 6.4)
	Unassisted	*vs.*	Deform. switches	0.7 * (0.1 to 1.3)	2.3 (-1.1 to 5.6)

Table 5.3: Bonferroni post-test multiple comparisons. The reported measures are mean difference, significance levels and (95% confidence intervals). Significance Levels are reported as * for ($0.01 < p \leq 0.05$), ** for ($0.001 < p \leq 0.01$), *** for ($0.0001 < p \leq 0.001$) and **** for ($p \leq 0.0001$).

5.4.3 Target shape

A one-way ANOVA was performed to test the null hypothesis that the target shape has a significant effect on the number of missed-clicks or experiment time. Reported results were the ANOVA p value, (F ratio and degrees of freedom [df]). The results do not show a statistically significant difference in the number of missed-clicks for target shape between deformable cones ($F_{2,28} = 0.33$, $p = 0.7185$) or deformable switches ($F_{2,28} = 2.63$, $p = 0.0895$). However, the shape had a significant effect on the experiment time for deformable cones ($F_{2,28} = 4.21$, $p = 0.0251$). Bonferroni-corrected post-test comparisons indicate that the mean ($\pm$ SD) experiment time for squares (23.3 ± 9.5s) was significantly longer than for wide rectangles (19.4 ± 6.3s), though did not significantly differ from either target shape for tall rectangles (21.1 ± 7.9s). The respective experiment times for deformable switches showed a similar pattern (24.7 ± 7.1s, 21.7 ± 4.6s, 22.3 ± 7.2s), though this did not reach statistical significance ($F_{2,28} = 2.96$, $p = 0.0685$).

Taken together, the results suggest that target shape has a less consistent and smaller effect on the participants' performance than the haptic condition or the target size. Specifically, there was no effect on the number of missed-clicks but participants took longer to click on squares than on rectangles. However, this effect was only significant for deformable cones between square and wide rectangular targets, with the former taking 3.9s longer to complete on average (95% CI 0.5-7.3s). Given that the rectangular shaped targets have twice the surface area of square targets it is unsurprising that a small difference has been observed. A future experiment designed specifically to analyse shape may provide a greater insight into the effect on user performance.

5.4.4 Discussion

Experiment 3 was designed to investigate the effect that target size and shape have on the performance of deformable cones and deformable switches. The main emphasis was on the benefits of providing haptic assistance for small sized targets given that previous studies had reported the most significant improvements under these conditions. The results show that the two deformable techniques enabled participants to select very small targets with a low error count. The level of significance for the number of missed-clicks and experiment time was shown to increase as the target size decreased. As computer monitors continue to grow in size and resolution it is likely that a larger gain will be required to navigate the whole of the screen. An increase in gain will reduce the effective width of the targets and make them difficult to select. The unassisted experiments indicate that pointing is most problematic when targets are small in the motor space. Haptic assistance could be very beneficial for motion-impaired computer users in the future. The results of the study suggest that the shape of the target has a less significant effect on participant performance.

Previous studies have also reported that haptic assistance provides the most significant improvements for people with more severe motion impairments [KLCR00] [HKL+01]. In this study the participant denoted by △ exhibited a tendency to slip when performing clicks. This is reflected in the click-release displacement and number of missed-clicks recorded during Experiment 1 for the unassisted interface. The most dramatic improvements for the participant denoted by △ were recorded for the small sized targets in Figure 5.15. The addition of deformable cones and deformable switches helped to prevent the cursor from leaving the target region, which significantly reduced error rates and improved targeting times. In contrast the motor skills of the participant denoted by □ tend to be more controlled, which means that the assistance has a less profound effect. However, the two deformable techniques were still able to reduce the number of missed-clicks and improve targeting times for this person, especially when the targets were small.

5.4.5 Conclusion

This section has presented a study investigating how target size and shape affect the performance of deformable cones and deformable switches for five motion-impaired participants. The ISO 9241-9 multidirectional point-and-click task was chosen as the interface to evaluate the assistance. The haptic techniques have shown to be most effective in extreme cases. The most significant improvements were recorded for small sized targets, which may be beneficial in the future as computer screens continue to increase in size and resolution. In addition, participants with less accurate clicking abilities also benefitted significantly from the assistance.

One of the advantages of the deformable techniques over traditional gravity wells is that they do not need to be tuned for different target sizes. For example, the spring constant of gravity wells would have to be adjusted for different sized targets due to the variation in displacement and thus force. In contrast, the operator can use the deformable haptic assistance straight away without the need of any interface specific calibration.

The results from this study show that the deformable haptic conditions are generalisable for different graphical user interfaces that contain targets with differing size and shape. It is anticipated that the assistance will be useful in providing access to existing computer software. The feature extractor discussed in Appendix B.2 has been designed to automatically generate haptic assistance for any Win32, Windows Forms or Windows Presentation Foundation (WPF) interface.

5.5 Experiment 4

5.5.1 Experimental procedure

Experiment 4 analyses the effectiveness of the haptic workbox that was presented in Section 3.4.8. The system has been designed to improve cursor navigation by allowing the operator to perform coarse movements with rate control and fine movements with direct positional control. The workbox also aims to reduce the effect of target distracters by disabling them when manipulating the cursor under rate control. Three separate experiments were conducted to evaluate the technique. Part 1 investigates the workbox in comparison to an unassisted interface and determines the effect that the magnification level has on interaction. Part 2 concentrates on the effectiveness of the workbox when used in conjunction with targeting haptic cues such as gravity wells. Finally, part 3 investigates the use of the workbox for people with severe motion impairments.

To evaluate the technique in a realistic environment the participants were asked to perform the on-screen keyboard task described in Section 4.2.2. Each operator was required to create a predefined sentence within a textbox using the OSK. The active target was always highlighted in red and only progressed to the next key once acquired. A sequence comprised of 50 successful target selections and data collection began after the first target had been selected. Data collection was continuous within a sequence and breaks of a minimum of 10 minutes were taken between trials. In accordance with Fitts' law participants were instructed to work as quickly as possible whilst maintaining a high level of accuracy. Due to only having a small sample size the experiments were repeated three times for each haptic condition. The order of presentation of the haptic techniques was randomised for each person. For the unassisted condition the force feedback provided by the haptic cues was turned off, meaning the Phantom Omni operated as an "ordinary" pointing device. Sessions were two hours long and shared between the user group.

The cursor movements of each person were recorded using the data recorder discussed in Section 4.2.5 and analysed according to the measures described previously in Sections 2.7 and 3.2. The main emphasis of the cursor analysis is on error rates and the effect that target distracters have on interaction. The participants' physical background and corresponding legend can be found in Table 4.2. Please note that the participant denoted by $\triangle$ was unavailable for this study.

5.5.2 Part 1

Experiment 4 part 1 is designed to show the effectiveness of the workbox in comparison to an unassisted interface. The study also investigates the effect that the magnification level has on performance. Three haptic conditions are investigated: unassisted, 20mm $\times$ 20mm sized workbox with one-to-one mapping and a 40mm $\times$ 40mm sized workbox with $\times$3 magnification. The gain levels remained the same for positional and rate control throughout the study.

The size of the workbox will determine the distance that the device has to travel over the course of the experiment. A small sized workbox of 20mm $\times$ 20mm was chosen for the one-to-one mapping to reduce the distance travelled under positional control. When the workbox is small in size and the magnification level is high then the targets will tend to dominate the interface. This is undesirable when under rate control because it requires more precise positioning of the black square to encapsulate the whole target. Figure 5.16 shows an example where a magnified target dominates the workspace. If the operator does not accurately place the black square around the target then only a percentage of it will be available to select when under positional control. When increasing the magnification level it is desirable to also increase the size of the workbox to ensure that it is large enough to incorporate the magnified targets. Therefore, the workbox size was increased to 40mm $\times$ 40mm for the $\times$3 magnification level.

The data has been described for each haptic condition using the mean and standard error of the mean (SEM), with data for each participant presented. The results are from five participants performing fifty successful selections with three repetitions.

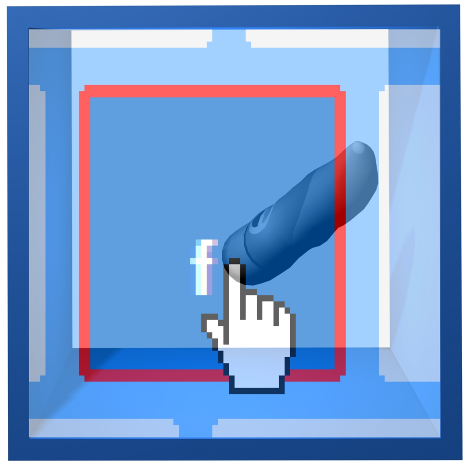

Figure 5.16: When the magnification level is high and the workbox is small in size, then the targets will dominate the workspace. This is undesirable when under rate control because it requires more precise positioning of the black square to encapsulate the whole target. The figure shows the Windows on-screen keyboard with the magnified f key dominating the workspace of the workbox.

For each cursor measure a reduction in magnitude is desirable when compared to the unassisted experiment and will signify an improvement. To test the null hypothesis that there were no differences between haptic conditions for each cursor measure, a repeated measures one-way ANOVA was used with planned post-test comparisons. The planned comparisons were as follows: unassisted against workbox one-to-one mapping, unassisted against workbox $\times 3$ magnification and workbox one-to-one mapping against workbox $\times 3$ magnification. To correct for multiple comparisons the method of Bonferroni was used. Reported results were the ANOVA p value, (F ratio and degrees of freedom [df]), difference in means between groups for each planned comparison with 95% confidence intervals (95% CI) for the difference between the two group means and statistical significance given multiple comparisons.

Missed-clicks

A one-way ANOVA shows that the haptic condition did not yield a significant effect on the number of missed-clicks ($F_{2,8} = 1.05$, $p = 0.3936$). The results are presented in Figure 5.17. Previous studies have reported an increase in error rates when using rate controlled devices and so it is important to allow direct positional control for accurate targeting [Epp86] [CEB87] [MD96]. No adverse effects were reported for the number of missed clicks for either workbox condition.

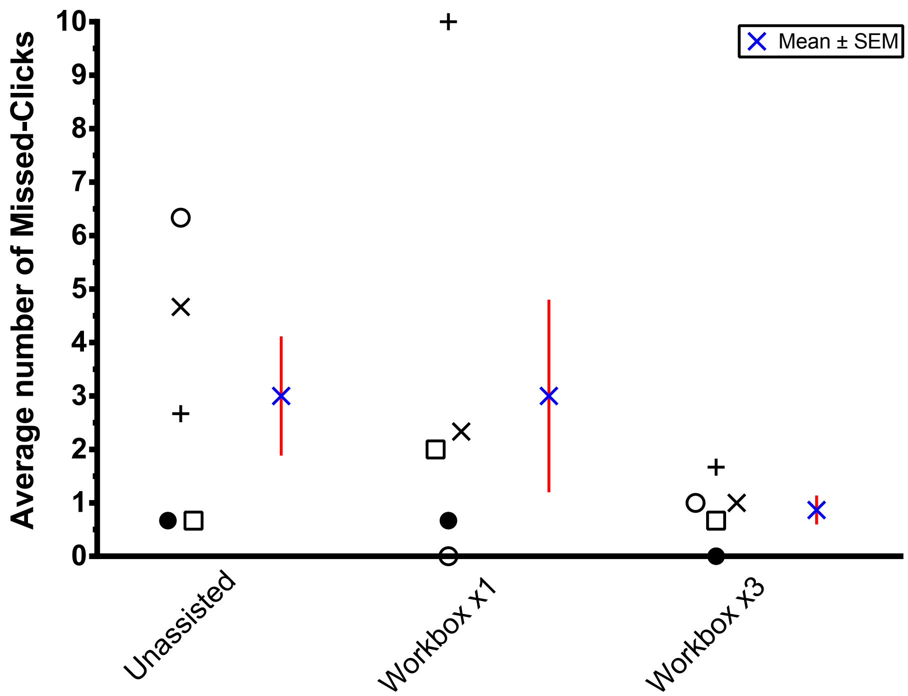

Figure 5.17: The number of missed-clicks recorded over fifty successful selections for each haptic condition with three repetitions.

The ×3 magnification workbox reduced the mean number of missed-clicks compared to the other conditions but did not reach statistical significance in Table 5.4. It is likely that the workbox magnification level would have a more significant effect when compared against smaller sized targets, as was shown in Experiment 3.

Experiment time

A one-way ANOVA shows that the haptic condition has a significant effect on the experiment time ($F_{2,8} = 30.7$, $p = 0.0002$). The results are presented in Figure 5.18. The planned comparisons in Table 5.4 show that both workbox conditions took significantly longer than the unassisted interface. Previous studies have shown that rate control is slower than direct positional control because involuntary tremor causes changes in the cursor velocity, which makes it difficult to stop precisely at a desired location on the screen [MD96]. The experiment time was significantly longer for the workbox with $\times 3$ magnification compared to the one-to-one mapping condition. This is unsurprising given that the interface is essentially $\times 3$ of the size and so it will take longer to navigate.

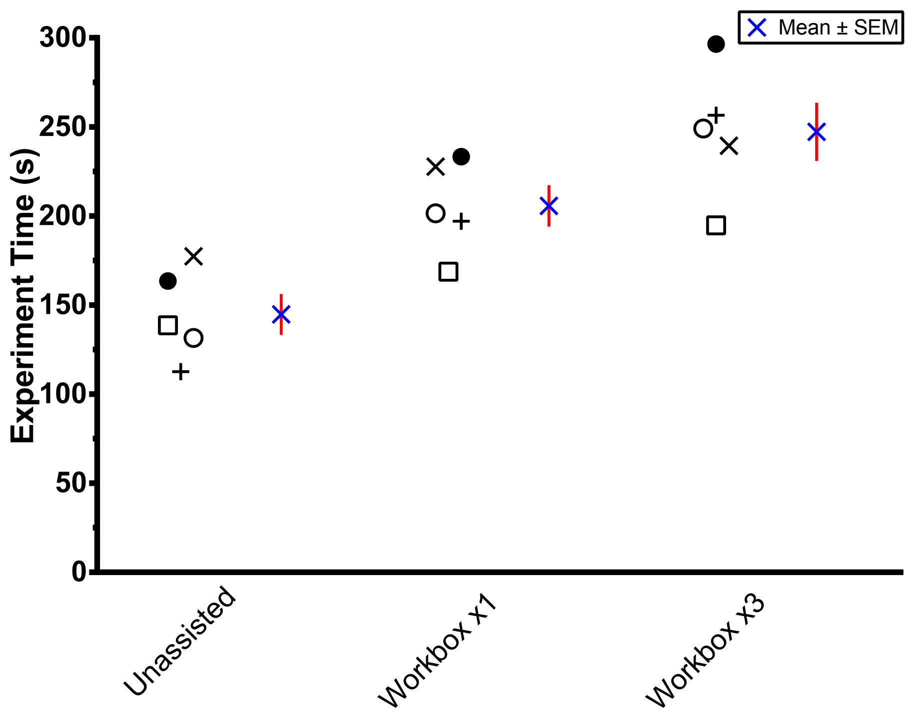

Figure 5.18: The experiment time recorded over fifty successful selections for each haptic condition with three repetitions.

Experiment distance travelled

A one-way ANOVA shows that the haptic condition has a significant effect on the experiment distance travelled ($F_{2,8} = 10.3$, $p = 0.0061$). The workbox was designed to reduce the distance travelled by utilising rate control for coarse navigation. However, the results in Figure 5.19 and Table 5.4 show that the one-to-one mapping workbox did not differ significantly compared to the unassisted interface. It is likely that the limitations of rate control meant that more corrections were required to position the black square in the desired location. The distance travelled for the ×3 magnification workbox was significantly greater than the one-to-one mapping condition and the unassisted interface. However, it should be noted that the rate control ensured that the overall distance travelled was less than ×3 of that recorded for the two other conditions.

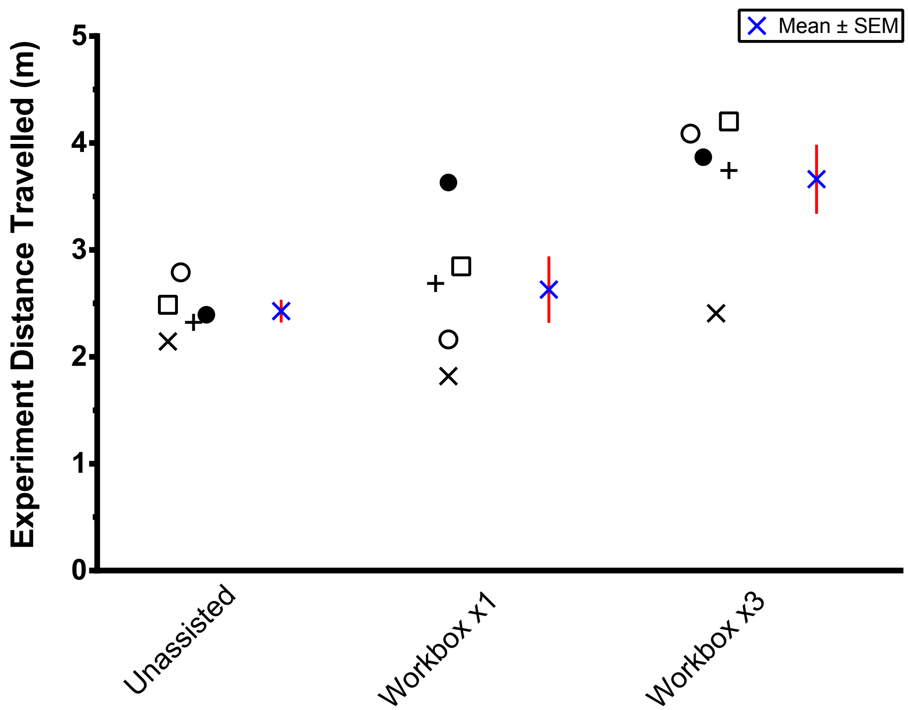

Figure 5.19: The experiment distance travelled recorded over fifty successful selections for each haptic condition with three repetitions.

Results of multiple comparisons

To provide information on the statistical significance of the differences a repeated measures one-way ANOVA was performed with planned post-test comparisons. Since the null hypothesis was rejected a Bonferroni multiple comparison was used to determine which means are different. The results are shown in Table 5.4.

Planned Comparisons			Missed-Clicks	Experiment Time (s)	Distance Travelled (m)
Unassisted	*vs.*	Workbox one-to-one	0.0 (-5.1 to 5.1)	-60.9 ** (-101 to -21.2)	-0.20 (-1.08 to 0.68)
Unassisted	*vs.*	Workbox $\times 3$	2.1 (-3.0 to 7.3)	-103 *** (-142 to -62.8)	-1.23 ** (-2.11 to -0.36)
Workbox one-to-one	*vs.*	Workbox $\times 3$	2.1 (-3.0 to 7.3)	-41.6 * (-81.3 to -1.91)	-1.03 * (-1.91 to -0.15)

Table 5.4: Bonferroni post-test multiple comparisons. The reported measures are mean difference, significance levels and (95% confidence intervals). Significance levels are reported as * for ($0.01 < p \leq 0.05$), ** for ($0.001 < p \leq 0.01$), *** for ($0.0001 < p \leq 0.001$) and **** for ($p \leq 0.0001$).

5.5.3 Part 2

Experiment 4 part 2 is designed to show the effectiveness of the workbox when used in conjunction with targeting haptic cues such as gravity wells. Gravity wells have been included in this study because they are the most widely reported haptic assistance and are transferable to 2DOF devices, as is the workbox itself. By introducing target distracters it will be possible to determine how effective the workbox is at reducing their intrusiveness. The size of the workbox will determine the number of distracters that lie within the proximity of the target when under positional control. A small sized workbox of 20mm $\times$ 20mm was chosen to reduce the number of potential distracters that could lie inside the workspace at any given time. Three haptic conditions are investigated: gravity wells, 20mm $\times$ 20mm sized workbox one-to-one mapping and 20mm $\times$ 20mm sized workbox one-to-one mapping with gravity wells.

The data has been described for each haptic condition using the mean and standard error of the mean (SEM), with data for each participant presented. The results are from five participants performing fifty successful selections with three repetitions. For each cursor measure (apart from experiment time spent on the virtual plane) a reduction in magnitude is desirable and will signify an improvement. To test the null hypothesis that there were no differences between haptic conditions for each cursor measure, a repeated measures one-way ANOVA was used with planned post-test comparisons. The planned comparisons were as follows: gravity wells against workbox, gravity wells against workbox with gravity wells and workbox against workbox with gravity wells. To correct for multiple comparisons the method of Bonferroni was used.

Reported results were the ANOVA p value, (F ratio and degrees of freedom [df]), difference in means between groups for each planned comparison with 95% confidence intervals (95% CI) for the difference between the two group means and statistical significance given multiple comparisons.

Missed-clicks

A one-way ANOVA shows that the haptic condition did not have a significant effect on the number of missed-clicks ($F_{2,8} = 1.69$, $p = 0.2441$). The results are presented in Figure 5.20. No adverse effects were reported for the workbox with gravity wells compared to the gravity well only condition. The results are consistent with those reported in Experiment 1, where gravity wells did not produce statistically significant improvements. However, the main concentration in this study is the effectiveness of the workbox for reducing the intrusiveness of target distracters.

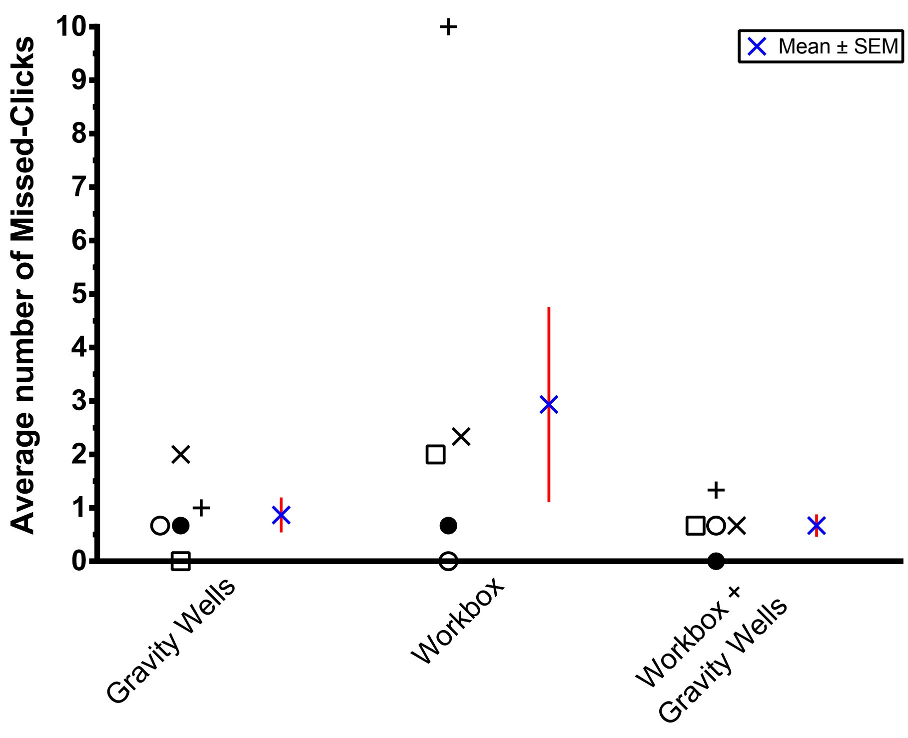

Figure 5.20: The number of missed-clicks recorded over fifty successful selections for each haptic condition with three repetitions.

Experiment time

A one-way ANOVA shows that the haptic condition has a significant effect on the experiment time ($F_{2,8} = 10.1$, $p = 0.0064$). The results are presented in Figure 5.21. During the previous studies presented in this thesis it was observed that the experiment time would often increase when target distracters are introduced that cannot be easily ignored such as gravity wells. The workbox was designed to reduce the effect of target distracters by disabling all haptic cues when under rate control. However, no significant differences in experiment time were recorded in Table 5.5 for the workbox condition compared to the workbox with gravity wells. The lack of improvement suggests that the majority of difficulties with target distracters lie primarily when under positional control rather than rate control. Once again the experiment time was significantly longer for the workbox conditions compared to the positionally controlled gravity well only condition.

Percentage of experiment time spent on the virtual plane

The amount of time spent in contact with the virtual plane will give an indication of the intrusiveness of target distracters when haptic cues are used in conjunction with the workbox. The workbox only condition will give a benchmark for the natural amount of time that participants spend in contact with the virtual plane. A one-way ANOVA shows that the haptic condition has a significant effect on the percentage of experiment time spent on the virtual plane ($F_{2,8} = 23.0$, $p = 0.0005$). The results are presented in Figure 5.22. The workbox was designed to reduce the effect of target distracters by disabling all haptic cues when under rate control. However, the results in Table 5.5 show that the participants spent significantly less time in contact with the virtual plane for the workbox with gravity wells condition compared to the workbox alone. Although there are benefits from disabling target distracters under rate control, it seems likely that the majority of issues occur from neighbouring haptic cues when under positional control. Once the operator experiences intrusion from target distracters they will then make a conscious effort to lift off the virtual plane to pass over them. The planned comparisons also show that participants spent significantly less time in contact with the virtual plane for the gravity well condition

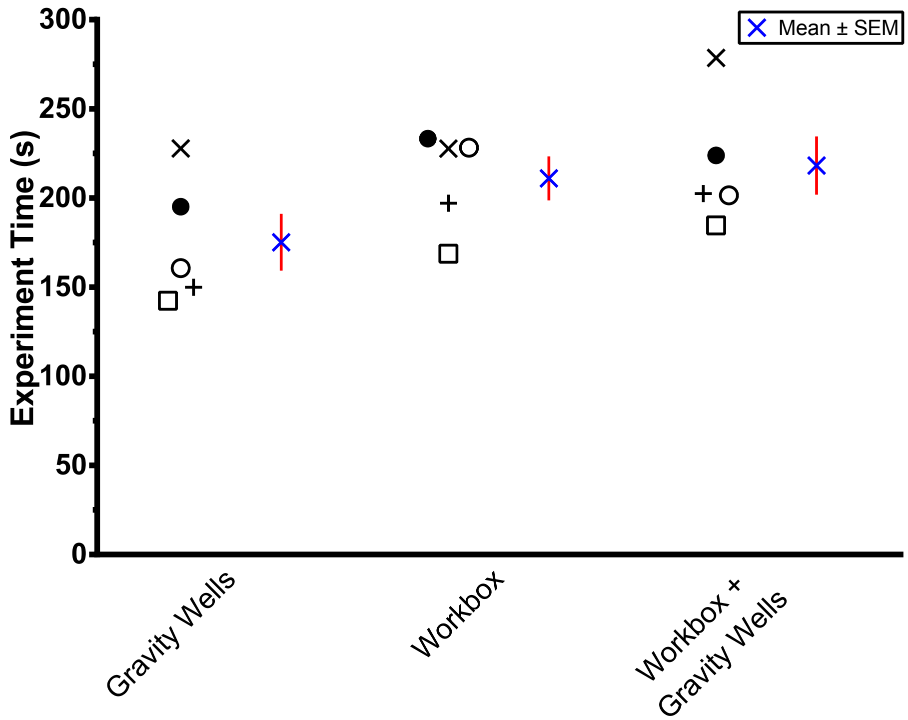

Figure 5.21: The experiment time recorded over fifty successful selections for each haptic condition with three repetitions.

compared to the workbox alone.

Experiment distance travelled

A one-way ANOVA shows that the haptic condition has a significant effect on the experiment distance travelled ($F_{2,8} = 16.4$, $p = 0.0015$). The results are presented in Figure 5.23. The previous experiments in this thesis have reported an increase in the experiment distance travelled when target distracters are introduced. This is due to the additional corrections required when exiting and overshooting distracters. The workbox was designed to reduce the distance travelled by utilising rate control for coarse navigation and temporarily disabling target distracters. However, the planned comparisons in Table 5.5 show that the distance travelled was significantly greater for the workbox with gravity wells compared to the workbox alone. Once again the results suggest that the majority of issues with distracters occur from neighbouring

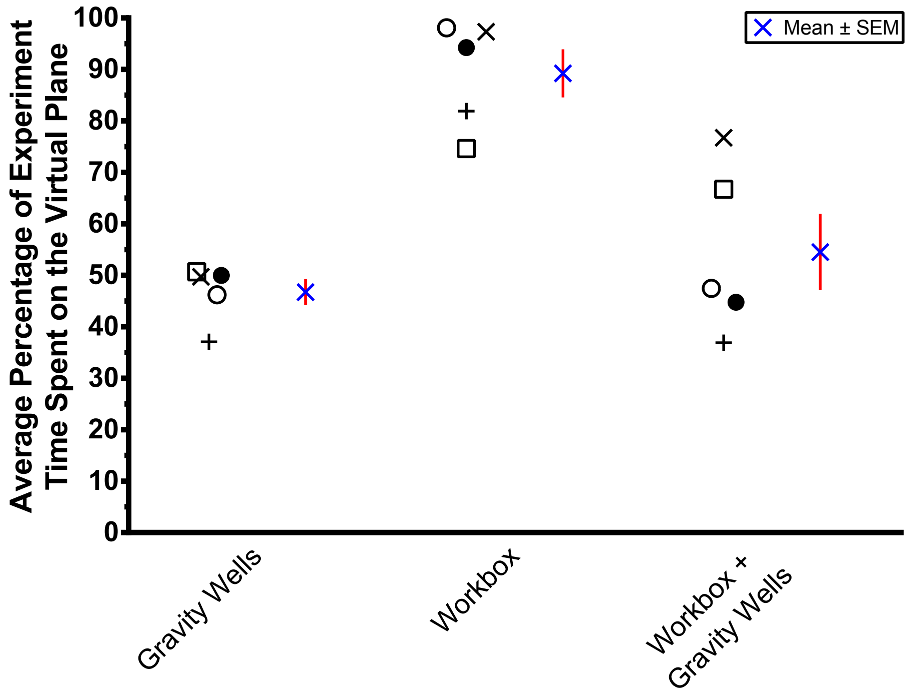

Figure 5.22: The percentage of experiment time spent on the virtual plane recorded over fifty successful selections for each haptic condition with three repetitions.

targets when under positional control.

Results of multiple comparisons

To provide information on the statistical significance of the differences a repeated measures one-way ANOVA was performed with planned post-test comparisons. Since the null hypothesis was rejected a Bonferroni multiple comparison was used to determine which means are different. The results are shown in Table 5.5.

Distracters

The results from the cursor analysis did not show the predicted improvements for the workbox in terms of reducing the effects of target distracters. The cursor traces were analysed using the playback feature described in Section 4.2.5 in an attempt to understand why this was the case. The workbox was designed to reduce the

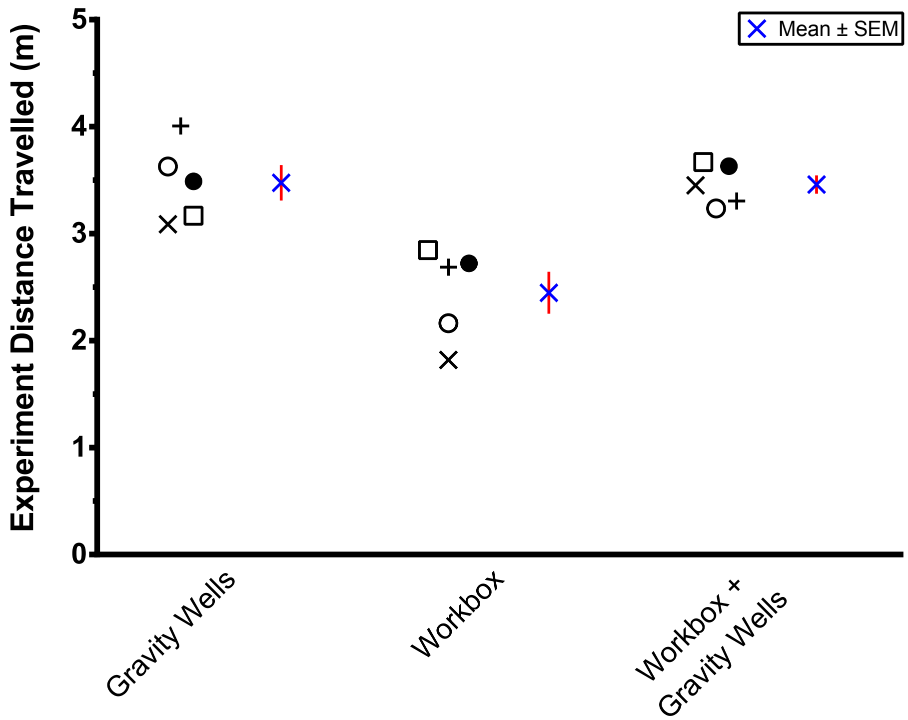

Figure 5.23: The experiment distance travelled recorded over fifty successful selections for each haptic condition with three repetitions.

effect of distracters by disabling the haptic cues when under rate control. Within the predefined sentence the longest transition across multiple distracters was from the P key to E. Figure 5.24 shows the cursor trajectories of two participants using gravity wells with and without the workbox. It is clear that there is considerably less positional disruption to the cursor through other targets along the task axis when the workbox is used. Similar cursor traces were observed for the other participants.

However, the benefits reported when under rate control were not as significant as the detrimental effect of target distracters when under positional control. The workbox size was purposely chosen to be small to limit the number of distracters that could lie inside the workbox at any given time. However, the cursor traces in Figure 5.25 indicate that the majority of difficulties with distracters lie primarily with neighbouring targets when under positional control. It is the haptic cues that are in close proximity of the target that cause positional disruption to the cursor that the

Planned Comparisons			Missed-Clicks	Experiment Time (s)	% Time on Virtual Plane	Distance Travelled (m)
Gravity Wells	*vs.*	Workbox one-to-one	-2.1 (-6.4 to 2.1)	-35.8 * (-66.7 to -4.9)	-42.5 *** (-62.6 to -22.4)	1.03 ** (0.41 to 1.65)
Gravity Wells	*vs.*	Gravity Wells & Workbox one-to-one	0.2 (-4.0 to 4.4)	-43.0 ** (-73.9 to -12.2)	-7.8 (-27.9 to 12.4)	0.017 (-0.60 to 0.64)
Workbox one-to-one	*vs.*	Gravity Wells & Workbox one-to-one	2.3 (-1.9 to 6.6)	-7.2 (-38.1 to 23.7)	34.7 ** (14.6 to 54.9)	-1.01 ** (-1.63 to -0.39)

Table 5.5: Bonferroni post-test multiple comparisons. The reported measures are mean difference, significance levels and (95% confidence intervals). Significance levels are reported as * for ($0.01 < p \leq 0.05$), ** for ($0.001 < p \leq 0.01$), *** for ($0.0001 < p \leq 0.001$) and **** for ($p \leq 0.0001$).

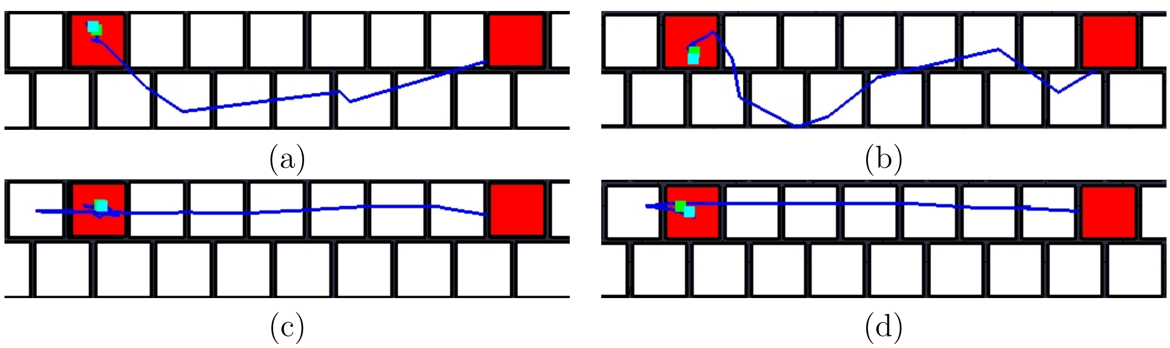

Figure 5.24: The cursor trace between keys P and E for the participants denoted by ●(a) and ×(b) using gravity wells. The cursor trace between keys P and E for the participants denoted by ●(c) and ×(d) using the workbox with gravity wells.

operator then has to correct for. This is consistent with the results reported in the cursor analysis.

5.5.4 Part 3

An additional participant has been included in the study to show the potential of the workbox for people with severe motion impairments. The participant denoted by ⋆ has cerebral palsy, which limits their range of movement and makes it difficult to grasp. The adjustable strap shown in Figure 5.26 was designed for the Phantom Omni to assist with grasping the stylus.

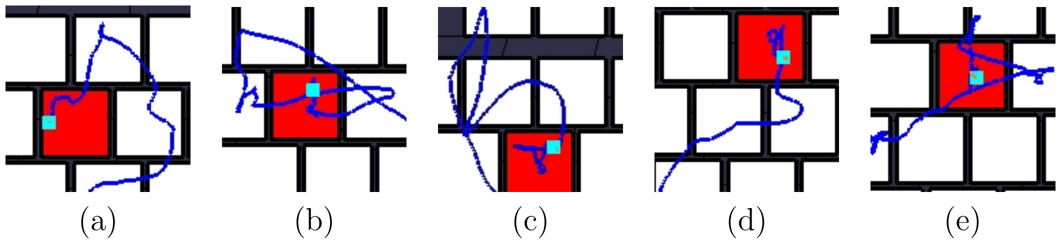

Figure 5.25: The cursor traces from the participants denoted by □(a), ●(b), ×(c), ○(d), +(e) show examples of the intrusion from target distracters when using the workbox with gravity wells under direct positional control.

Figure 5.26: The adjustable strap designed for the Phantom Omni to assist with grasping the stylus.

The participant denoted by ⋆ is able to navigate an electric wheelchair using a displacement joystick. Their usual computer input method is the "Camera Mouse", which was discussed in Section 2.4.2 [BGF02]. The participant is unable to operate the device switch and so it was not possible to include them in the point-and-click task on this occasion. However, to demonstrate the potential benefits of the workbox for someone with a limited range of movement the participant denoted by ⋆ was asked to position the cursor within four squares at the extremities of the screen using unassisted direct mapping, the workbox and the "Camera Mouse". The cursor trajectories of the three attempts are shown in Figure 5.27. The range of movement that the participant was able to produce in the unassisted experiment was approximately 80mm × 80mm in device displacement. Cursor control at the extremities of the operator's range of movement was less controlled and so a workbox size of 40mm $\times$ 40mm was chosen. The workbox was positioned so that its origin was centred around the operator's range of movement.

The cursor trace in Figure 5.27(a) shows that the participant denoted by $\star$ was unable to reach the four squares using the Phantom Omni with direct positional control. However, the cursor traces for the workbox in Figure 5.27(b) and the "Camera Mouse" in Figure 5.27(c) show that the assistance enabled the participant to reach all of the targets. The cursor trace for the workbox is visibly more controlled than the "Camera Mouse". The paths recorded for the workbox are much more direct and there are significantly fewer corrections. One of the limitations of the "Camera Mouse" is that it is very sensitive to user input, which can make precise manipulations difficult to perform. The participant denoted by $\star$ regularly plays a colour pairing game at the Norfolk and Norwich Scope Association (NANSA). The game does not contain accessibility features, which means that the interface extractor discussed in Appendix B.2 is ineffective. As a result, a map generator has been created so that haptic assistance can be manually positioned for non-supported interfaces. Further details are presented in Appendix B.3.

5.5.5 Discussion

The workbox technique evaluated in Experiment 4 was designed to allow the operator to rapidly navigate large and complex GUIs with a haptic feedback device, whilst still permitting accurate target selection. As screen size and resolutions continue to rise this will become more important especially for haptic devices with a limited workspace.

Experiment 4 part 1 showed that the decision to provide direct positional control for clicking was justified given that there were no significant differences reported between conditions for the number of missed-clicks. However, it appears that the rate control has a detrimental effect on the experiment time, which is similar to previous studies that have investigated rate controlled joysticks. The inclusion of rate control was designed to reduce the distance travelled by the operator but the results for the one-to-one mapping workbox did not differ significantly compared to the unassisted

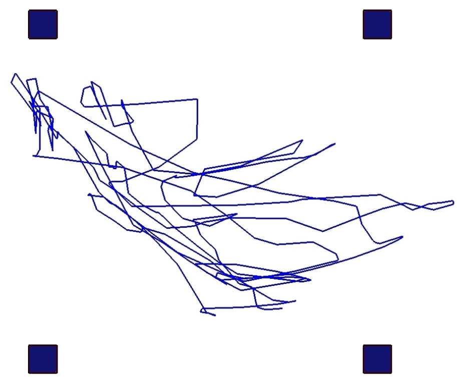

(a)

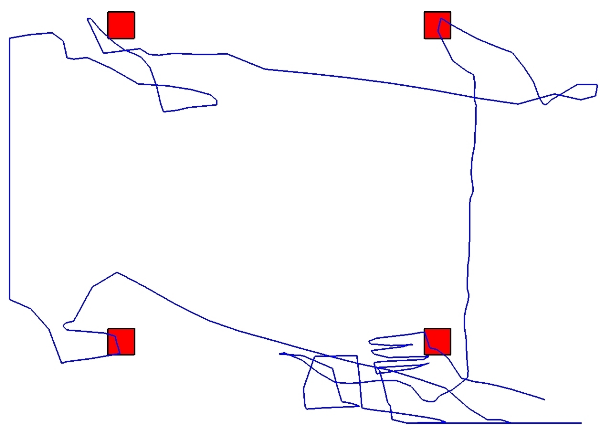

(b)

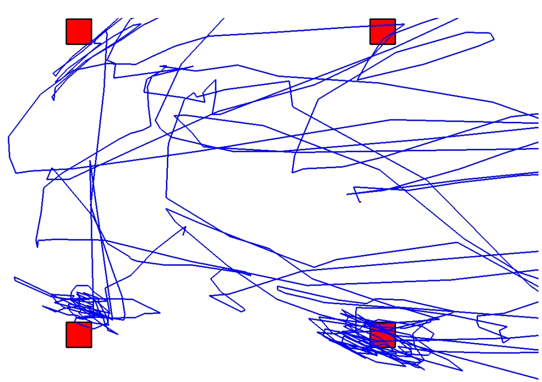

(c)

Figure 5.27: The participant denoted by $\star$ was asked to place the cursor within four targets at the extremities of the screen. The respective cursor traces are: unassisted direct mapping with the Phantom Omni (a), 40mm $\times$ 40mm workbox with the Phantom Omni (b) and the "Camera Mouse" (c).

interface. It is likely that the limitations of rate control meant that more corrections were required to position the black square in the desired location. The ×3 magnification level reduced the mean number of missed-clicks but did not reach statistical significance in the planned comparisons. It is likely that the workbox magnification level would have a more significant effect when compared against smaller sized targets.

Experiment 4 part 2 showed that the workbox was ineffective at reducing the effect of target distracters when gravity wells are introduced. The workbox was designed to reduce the intrusiveness of target distracters by disabling them when under rate control and limiting the number inside the workbox when under positional control. However, the results show that participants spent significantly less time in contact with the virtual plane for the workbox with gravity wells compared to the workbox alone. Once the operator experiences intrusion from target distracters they will then make a conscious effort to lift off the virtual plane to avoid them. The planned comparisons show that the distance travelled was significantly greater for the workbox with gravity wells compared to the workbox alone. The increase in distance travelled is caused by the subsequent corrections that are required when passing through target distracters. These results suggest that the majority of issues with distracters occur when under positional control. The cursor traces were analysed and confirm that the majority of difficulties arise from neighbouring haptic cues when using direct positional control.

Experiment 4 part 3 showed the potential benefits of the workbox for people with severe motion impairments. A participant with only a limited range of movement (80mm × 80mm) was able to navigate the cursor accurately across the whole screen using a 40mm × 40mm workbox. The cursor trace was significantly more controlled for the workbox compared to the operator's usual input method of the "Camera Mouse".

5.5.6 Conclusion

This section has presented a study investigating the performance of the workbox for motion-impaired participants. The results produced from the analysis of the workbox appear to be mixed. There were no significant differences reported in the number of missed-clicks but participants took significantly longer to perform the tasks when using the workbox. The overall lack of improvement could be because the majority of difficulties lie primarily in clicking rather than in navigating to the target [LHK+02b]. The secondary purpose of the workbox was to reduce the distance travelled and limit the effect of target distracters. However, the results suggest that the technique was ineffective in both of these areas.

The greatest potential of the workbox was reported for people with more severe motion impairments. The technique was shown to improve computer access for a person with a limited range of movement by enabling them to navigate all of the computer screen with a combination of rate and position control. Since the development of deformable switches the participant denoted by $\star$ has been able to play the puzzle game described in Appendix B.3 using the Phantom Omni with the workbox.

Chapter 6

Conclusions

6.1 Introduction

In the most recent decades researchers from a diverse set of disciplines have been investigating methods to improve computer access for people with motion impairments. Suitable access to a computer interface provides people with much greater opportunities in their personal and professional development. However, most user interfaces are designed to meet the requirements of the mass market, which means they are often inaccessible to people with physical disabilities. As a result, it is more useful to adapt existing user interfaces to improve access to commonly used software. Early research into haptic assistance showed its potential under experimental conditions but also highlighted a number of limitations for real world applications. The haptic assistance proposed in this thesis has been designed specifically for motion-impaired participants. The focus has been on the development of 3DOF techniques to allow non-intrusive interaction with existing graphical user interfaces (GUIs). The results presented in this thesis have shown that the ability to integrate haptic feedback with existing software can significantly improve interaction rates. A conscious effort has been made to provide a simpler access technology that has a greater chance of long term adoption.

The following section discusses the main contributions of each of the methods investigated. These details are presented in the order they appear in this thesis. Section 6.4 outlines the future work that could follow this research.

6.2 Discussion

Chapter 2 discussed the difficulties that motion-impaired people experience with human-computer interaction (HCI). It is important as a researcher to have an understanding of these difficulties so that areas of assistance can be identified. A number of alternative input methods were reviewed to highlight their benefits and limitations. Previously investigated haptic assistance was reviewed and a number of shortcomings were highlighted including: device calibration, target distracters, the intrusive nature of traditional haptic techniques, the haptic trade-off and 2DOF devices. A series of guidelines have been proposed to give an insight into how haptic techniques should be designed to limit intrusion. The research from this chapter highlighted several significant challenges in the field. Therefore, three significant challenges related to the design and integration of haptic assistance with graphical user interfaces were investigated throughout this thesis:

1. Producing techniques that do not require force calibration for optimisation.
2. Alleviating the effects of target distracters.
3. Developing non-intrusive techniques that can be easily used or ignored.

Chapter 3 discussed the implementation of the haptic assistance evaluated in this thesis. A number of new haptic assistive techniques were proposed that have been designed to overcome the shortcomings identified in Chapter 2. The proposed techniques include: haptic cones, V-shaped funnels, deformable cones, virtual switches, deformable switches and the haptic workbox. A new haptic rendering algorithm was implemented to permit the development of deformable cones and deformable switches.

The experimental setup was presented in Chapter 4. The task often conducted with cursor analysis techniques is based on the ISO 9241-9 standard for pointing device evaluation. The ISO 9241-9 task has limitations when analysing haptic assistance because it does not take into consideration the effect of target distracters. In most GUIs the toolbar buttons are arranged in rows or columns and so the evenly spaced, circular layout is unrealistic. In this study the Windows on-screen-keyboard (OSK) was chosen as the primary interface to evaluate the haptic assistive techniques.

The densely populated GUI provides a more realistic and extensive evaluation of the effects of target distracters. In this study seven participants with varying degrees of motion impairment were recruited from the Norfolk and Norwich Scope Association (NANSA). A brief summary of the participants' physical background has been provided.

The results of the study were presented in Chapter 5. Experiment 1 investigated the performance benefits of haptic assistance that does not require force calibration compared to traditional techniques in a static configuration. Haptic cones and V-shaped funnels were compared against gravity wells and high-friction targets. The results show that haptic cones outperformed gravity wells and high-friction targets in terms of reducing error rates during a point-and-click task. Haptic cones reduced the mean number of missed-clicks by 71% and were the only technique to show a statistically significant improvement over the unassisted interface. The clamping at the cone apex helps maintain cursor stability when clicking. One of the major advantages of haptic cones is that they do not require force calibration to optimise interaction. As a result, the operator can use the assistance immediately and any changes in their clicking characteristics, such as fatigue, will not require force recalibration. The second major advantage of haptic cones and V-shaped funnels is that the operator does not have to oppose a force to exit target distracters.

Experiment 2 explored the benefits of non-intrusive haptic assistance that the operator can choose to use or ignore. The study concentrated on the evaluation of deformable haptic cones and deformable virtual switches. The analysis was focused on interaction rates and the effect that target distracters have on user performance in a densely populated interface. The results showed that the two deformable techniques were the most effective at improving targeting times and reducing error rates during a realistic point-and-click task. Deformable cones and deformable switches reduced the mean number of missed-clicks by 75% and 92% respectively in comparison to the unassisted experiment. The improvement in clicking accuracy will be credited to the effective clamping at the apex, which helps maintain cursor stability. Deformable haptic cones and deformable virtual switches both significantly improved the experiment time by 27% and 25% respectively in comparison to the unassisted interface.

Both deformable techniques were significantly more effective at reducing the experiment time in comparison to traditional gravity wells. This will be credited to the increased confidence during target selection without the intrusion from neighbouring haptic cues.

The cursor measures confirm that target distracters are less intrusive for deformable cones and deformable switches. The deformable techniques allow the operator to navigate more freely without imposing forces from neighbouring haptic cues. This is confirmed by the fact that there were no significant differences in the percentage of experiment time spent on the virtual plane when compared to the unassisted experiment. Many of the shortcomings highlighted in Section 2.6 have been alleviated by utilising a 3DOF interface to produce assistance that can be easily used or ignored. The ability to perform a gesture that enables the assistance will allow haptic feedback to be integrated with existing graphical user interfaces without intrusion from target distracters. This is important because it means that distracters are no longer the limiting factor in the development of haptic assistance.

Experiment 3 used the ISO 9241-9 task to investigate how target size and shape effect the performance of the deformable haptic assistance. The study was designed to show that the techniques are generalisable for real-world GUIs that contain different target sizes and shape. The results show that the two deformable techniques enabled participants to select very small targets with a low error count. The level of significance for the number of missed-clicks and experiment time increased as the target size decreased. The unassisted experiments indicate that pointing is most problematic when targets are small in the motor space. Given that target selection is the most common difficulty for motion-impaired computer users, it is encouraging that the assistance is so effective for small sized targets. As computer monitors continue to grow in size and resolution it is likely that the effective width of the targets will shrink, which will make them more difficult to select. Haptic assistance could be very beneficial to motion-impaired computer users in this area. The results of the study suggest that the shape of the target has a less significant effect on performance.

The haptic workbox investigated in Experiment 4 was designed to allow the operator to rapidly navigate large and complex GUIs with a haptic feedback device, whilst still permitting accurate target selection. This was achieved by creating a rate/position hybrid system. However, the results did not show significant improvements for targeting and the experiments were significantly slower when the workbox was used. The technique was ineffective at reducing the distance travelled and did not limit the intrusiveness of target distracters. The greatest potential of the workbox was reported for a person with a limited range of movement. The technique enabled the participant to navigate the whole of the computer screen with a combination of rate and position control.

6.3 Conclusions

This thesis explored haptic assistance as a means of improving computer access and interaction rates for people with physical impairments. The goal was to simplify and improve cursor efficiency. The main concentration has been on target acquisition techniques given that the literature reported that the majority of difficulties lie primarily in clicking rather than in navigating to the target. The new haptic techniques presented in this thesis utilise the three-dimensional attributes of the Phantom Omni to produce assistance that overcomes many of the limitations of previous 2DOF interfaces. The techniques were designed specifically for real world applications and were tested under extreme conditions using the densely populated Windows on-screen keyboard (OSK). Deformable cones and deformable switches were the most effective techniques for decreasing the number of missed-clicks and improving targeting times, the two measures of performance that are most important. The newly proposed cursor measures were useful in evaluating the performance of the haptic assistance and identifying the effects that target distracters have on interaction. Previously, target distracters have plagued user interfaces and were the limiting factor in the integration of haptic assistance for real world applications. The deformable techniques do not suffer from this shortcoming because the operator can choose when to use or ignore the assistance. The additional benefit of the techniques presented in this thesis is that they do not require force calibration, which allows the assistance to be used straight away. The haptic feedback is not intrusive because no forces are imposed on the operator, which means that the assistance will be better suited for a wider range of disabilities.

The results presented in this thesis were consistent with previous studies that reported the most significant improvements in extreme conditions. This is in reference to people with more severe motion impairments and when interface targets are small. In this study the participant denoted by $\triangle$ exhibited a tendency to slip when performing clicks. This was reflected in the click-release displacement and number of missed-clicks recorded for the unassisted interface. The addition of deformable cones and deformable switches helped to prevent the cursor from leaving the target region, which significantly reduced error rates and improved targeting times. In contrast the motor skills of the participant denoted by $\square$ tended to be more controlled, which meant that the assistance had a less profound effect. However, the two deformable techniques were still able to reduce the number of missed-clicks and improve targeting times for this person, especially when the targets were small. The most dramatic improvements for the participants as a whole were recorded for the small sized targets in Experiment 3. This confirms that the techniques presented in this thesis are generalisable for varying abilities and different interfaces. It is anticipated that the results will be useful in providing assistance that could significantly improve computer access.

To ensure that the haptic assistance is adopted it needs to be easily applied to existing software. The final interface presented in Appendix B has been designed to ensure that motion-impaired computer users can easily choose the haptic condition they wish to use and the application they wish to apply it to. The feature extractor obtains the shape and position of the interface buttons so that haptic assistance can be automatically generated for any Windows Win32, Windows Forms or WPF application. This allows the benefits of the techniques to be integrated with commonly used software such as the on-screen keyboard, Microsoft Office, Internet Explorer, etc. The map generator allows haptic maps to be manually created and saved for any non-Win32, Windows Forms or WPF interface. This will provide access to other applications such as computer games and puzzles. By overlaying the partially transparent interface window on top of the main OpenGL viewport it has been possible to provide suitable visual cues to accompany the haptic conditions. It is anticipated that the haptic assistance combined with existing software will provide people with much greater opportunities in their personal and professional development. The system has been adopted by many of the participants at the Norfolk and Norwich Scope Association (NANSA). The ability to interact with a computer more effectively has encouraged participants to use the haptic devices available at NANSA. The feature extractor has enabled commonly used software to be more easily accessible through the use of automatically generated haptic assistance. Companies that specialise in assistive technology such as Sensory Software may benefit from integrating haptic assistance with their toolkit. The Grid 2 system developed by Sensory Software is designed to improve communication and computer access through partitioning the interface into grid sets. The content of each grid is fully customisable to meet the communication and computer access needs of the user. A series of grids will normally be linked together such that selecting a particular cell on one grid will cause a different grid to be displayed. In this way large sets of information can be broken down into more easily manageable chunks for presentation to the user. Haptic assistance could be hugely beneficial to target selection in this type of environment.

As the age of the global population continues to rise, there is an increased demand for a more efficient and effective means of data entry into a computer. The work presented in this thesis has focused on improving computer access for people with motion impairments. However, the haptic assistance could be effective at improving interaction rates for able-bodied people. A number of studies have shown that advanced age can make cursor movement increasingly inaccurate. The general population is growing older and it is estimated that by 2020 almost half the adult population in the United Kingdom will be over 50, with the over 80's being the most rapidly growing sector. As computer usage spreads throughout the population, computer interfaces will have to adapt to meet the needs of the user group. Haptic interaction could have a major influence on this market especially as force feedback devices become more popular and affordable.

6.4 Future work

There are two main areas that could form the continuation of the work presented in this thesis. These are discussed in the following subsections.

6.4.1 Phantom Omni mouse

The most common pointing device used with a computer is the mouse. Many people with physical disabilities find the mouse difficult to operate but they often persevere with it because they are familiar with its operation. Trewin and Pain report that people often prefer standard mice or trackballs to specialised devices because of familiarity, availability and ubiquity [TP99]. This is consistent with other findings that report high abandonment and low adoption rates for alternative input methods, even amongst people with disabilities [RRW00] [Koe03].

If the Phantom Omni could be adapted so that it was used in a similar way to a traditional mouse then this would negate the initial training phase that is required for a new pointing device. The operator would still be able to utilise the 3DOF benefits of the interface such as the ability to lift the mouse off the desk to pass over target distracters. This is not possible with traditional 2DOF haptic devices because they are often fixed to a platform. The mouse grip presented in Figure 6.1 has been adapted from a traditional corded mouse. The existing switches are wired to a 6mm stereo socket that has been mounted inside the casing.

By providing a mouse based interface there is more flexibility to adjust the gain because the translation of the cursor is not limited to the device workspace. Ideally, the operator would be presented with a virtual plane just above the surface of the desk so that the benefits of the 3DOF techniques proposed in this thesis could be utilised. However, the Phantom Omni does not provide torques, which means that the device would not be able to support itself about the x-axis or the z-axis. It is also undesirable to load the motors for extended periods of time. A possible solution to these issues would be to integrate a sprung base within the mouse, which would reduce the workload on the motors and still allow access to the 3DOF capabilities of the interface. Further research is required to develop the hardware and to evaluate

Figure 6.1: The concept of using the Phantom Omni in a mouse configuration.

the performance of the haptic assistance in this configuration.

6.4.2 Free skate

One of the limitations of the techniques presented in this thesis is that they require knowledge of the location of the targets. Wobbrock et al. discuss the practicalities of methods that are target-aware compared to those that are target-agnostic [WFL09]. The feature extractor presented in Appendix B has been designed to automate any Windows Win32, Windows Forms or WPF application with haptic assistance. However, there may be occasions where the feature extractor may not be useful because the operator may wish to accurately click in a location that does not contain an interface button. For example, drawing a line in Microsoft Paint or moving a vertex in 3D Studio Max, etc.

To overcome these issues a method called "free skate" has been developed that will provide haptic clicking assistance at any location on the computer screen. The deformable techniques presented in this thesis have been adapted to follow the cursor so that the user can position the assistance at a location of their choice. Essentially, a single deformable cone or deformable switch is embedded into the virtual plane and translated with the on-screen cursor. The location of the assistance is locked in place when the operator deforms the cone/switch by 50% of its maximum depth. The user can then accurately click in the desired location and is unlikely to slip off the target. Once the cone/switch is 75% reformed then it will resume shadowing the cursor. The size of the assistance can be adjusted to meet the needs of the interface. The current implementation uses the modal button size within the chosen application.

The only potential disadvantage of the "free skate" method is that the operator will receive less assistance on-click compared to an interface that already has haptic cues located over the targets. Its success is reliant on the operator positioning the assistance accurately before locking it into position. However, the results presented in this study have shown that motion-impaired participants are more likely to slip off a target and so the assistance on-click may not be as essential. Further experiments need to be conducted to evaluate the performance of "free skate" compared to the static implementation and an unassisted interface. Qualitative feedback would be useful to determine if it is an intuitive method of interaction and to ensure that there are no intrusive limitations. The "free skate" feature has been made available in the final interface so as to provide assistance at locations that do not already contain haptic cues.

Appendix A

Demonstrations

The accompanying DVD contains video files captured during the running of the programs described in this thesis. The demonstrations are listed in the order in which they appear in the thesis. Each movie is captured with a standard video camera to illustrate both the screen and the user manipulating the haptic device.

A.1 Video playlist

1. Gravity wells
2. Gravity wells - target distracters
3. High-friction targets
4. High-friction targets - target distracters
5. Haptic cones
6. Haptic cones - target distracters
7. Haptic funnels
8. Haptic funnels - target distracters
9. Deformable cones
10. Deformable cones - target distracters
11. Virtual switches

12. Virtual switches - target distracters

13. Deformable switches

14. Deformable switches - target distracters

15. Haptic workbox

Appendix B

Final haptic interface

The following section discusses the implementation of the final system that has been developed to allow quick and easy access to haptic assistance for existing software. To ensure that the haptic assistance is adopted it needs to be accessible for users to operate independently. One of the shortcomings and major criticisms of alternative assistive technology is the need for complicated calibration and the amount of setting up time. The literature provided in Chapter 2 identified a number of guidelines that have been taken into consideration to avoid difficulties with the interface design and pointing device operations.

B.1 Graphical user interface (GUI) design

The design of the GUI will have a direct bearing on the system's usability. The interface has been configured so that motion-impaired computer users can choose their own assistance and the application they wish to apply it to. Screen size and resolution has increased significantly in recent years, which means that there is space available for docking additional assistive features. Figure B.1 shows the final GUI layout, which contains four OpenGL viewports that provide visual feedback and a toolbar that is used to configure the interface. The operator can choose between deformable cones and deformable switches by clicking on the corresponding buttons within the toolbar. The last assistive method that was employed is enabled by default. The operator can choose the software they wish to use by clicking the "Load Application" button and selecting the process from a list. The application window is then automatically docked on top of the main OpenGL viewport and automated with haptic assistance. Appendix B.2 describes the methods that have been employed for extracting features from the interface. By overlaying the partially transparent window on top of the main OpenGL viewport it has been possible to provide suitable visual cues to accompany the haptic conditions. Separate buttons have been provided in the toolbar to simulate the functionality of more difficult clicking operations such as right click, double-click and drag-and-drop. The drag-and-drop feature adopts the "free skate" method that is described in greater depth in Section 6.4.2. Buttons within the interface have been made as large as possible to ensure that they are easily accessible. Haptic assistance has also been provided for the features within the GUI. If the application is minimised or loses focus then the haptic cues are temporarily disabled until the window is restored. The advanced tab contains facilities that were used to conduct the experiments described in this thesis.

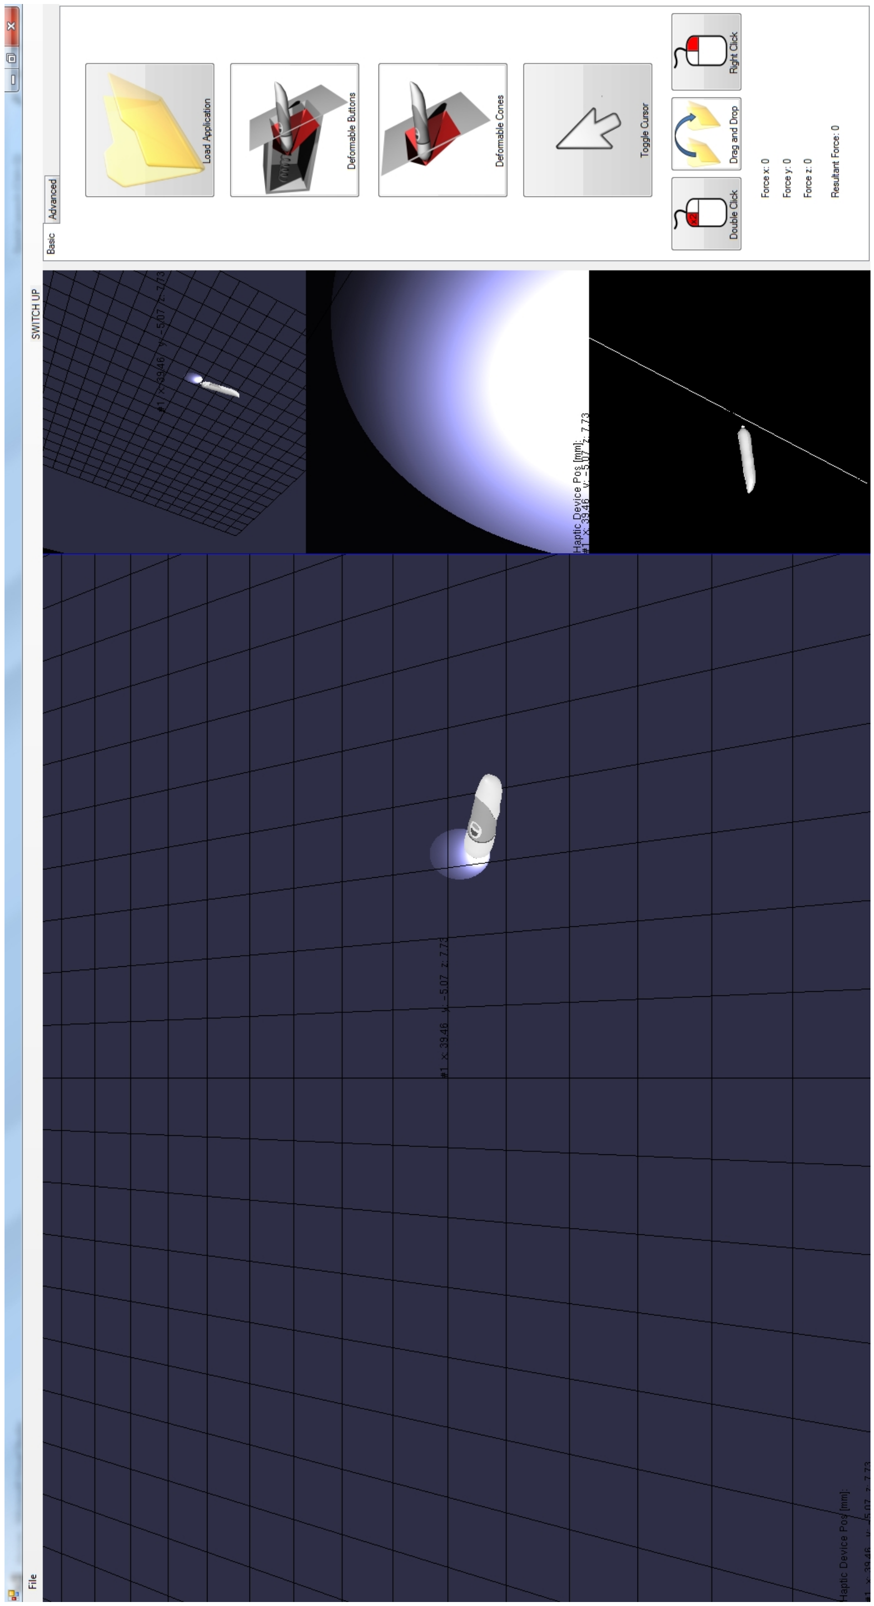

Figure B.1: The final graphical user interface. The operator can choose the haptic assistance they wish to use and the application they wish to apply it to. Separate buttons for right click, double-click and drag-and-drop operations have been provided.

B.2 Interface feature extraction

Programmatic access to existing graphical user interfaces (GUIs) is a crucial element for improving accessibility. The feature extractor presented in this section obtains the shape and position of the interface buttons so that haptic assistance can be automatically generated for any Win32, Windows Forms or Windows Presentation Foundation (WPF) application. This allows the benefits of the haptic assistance to be integrated with commonly used software such as the on-screen keyboard, Microsoft Office, Internet Explorer, etc.

The interface feature extraction has been achieved using the Windows UI Automation API. The new accessibility model provides programmatic access to information about the user interface. Developers of accessibility tools can use this information to create software that makes applications running on Windows more accessible to people with vision, hearing or motion impairments. The UI Automation API supplies the tools that are required for provider and client development. UI Automation providers are applications such as Microsoft Word, Excel, Internet Explorer or third-party software designed for the Windows operating system. UI Automation clients are assistive technology applications such as screen readers, screen magnifiers, alternative input methods, etc.

The UI Automation tree represents the entire user interface, where the root element is the current "Desktop" and child elements are application windows. The GUI elements are represented as nodes, which are children of the application to which they belong. In the UI Automation framework each `AutomationElement` exposes common properties of the user interface regardless of the underlying implementation (Win32, Windows Forms or WPF). One of these properties determined by the UI Automation provider is the `ControlType`, which describes the basic appearance and functionality of the control. For example, Button, ComboBox, CheckBox, RadioButton, ListBox, etc. Providers of controls for Win32, Windows Forms and WPF applications are supplied as part of the operating system. Custom providers can be developed for other UI frameworks or custom controls.

Inspect (Inspect.exe) is a Windows-based tool that enables the user to select any UI element and view the accessibility data. The navigational structure of the automation elements can be viewed in the UI Automation tree. Figure B.2 shows Inspect reporting the accessible information from a Microsoft Word application. The hierarchical tree of UI elements is shown on the left side. The main window element is the parent of a “MsoDockTop pane”, which is a parent of the highlighted “Bold” button. In the focus tracking mode, Inspect follows the UI item that has keyboard focus and automatically shows its properties in the tree view. Common uses for focus tracking include stepping through UI items to ensure that they can receive keyboard focus. In the hovering mode, Inspect reports the control properties when the mouse passes over an interface feature. Hover mode is more convenient when the user wants to select a specific UI item.

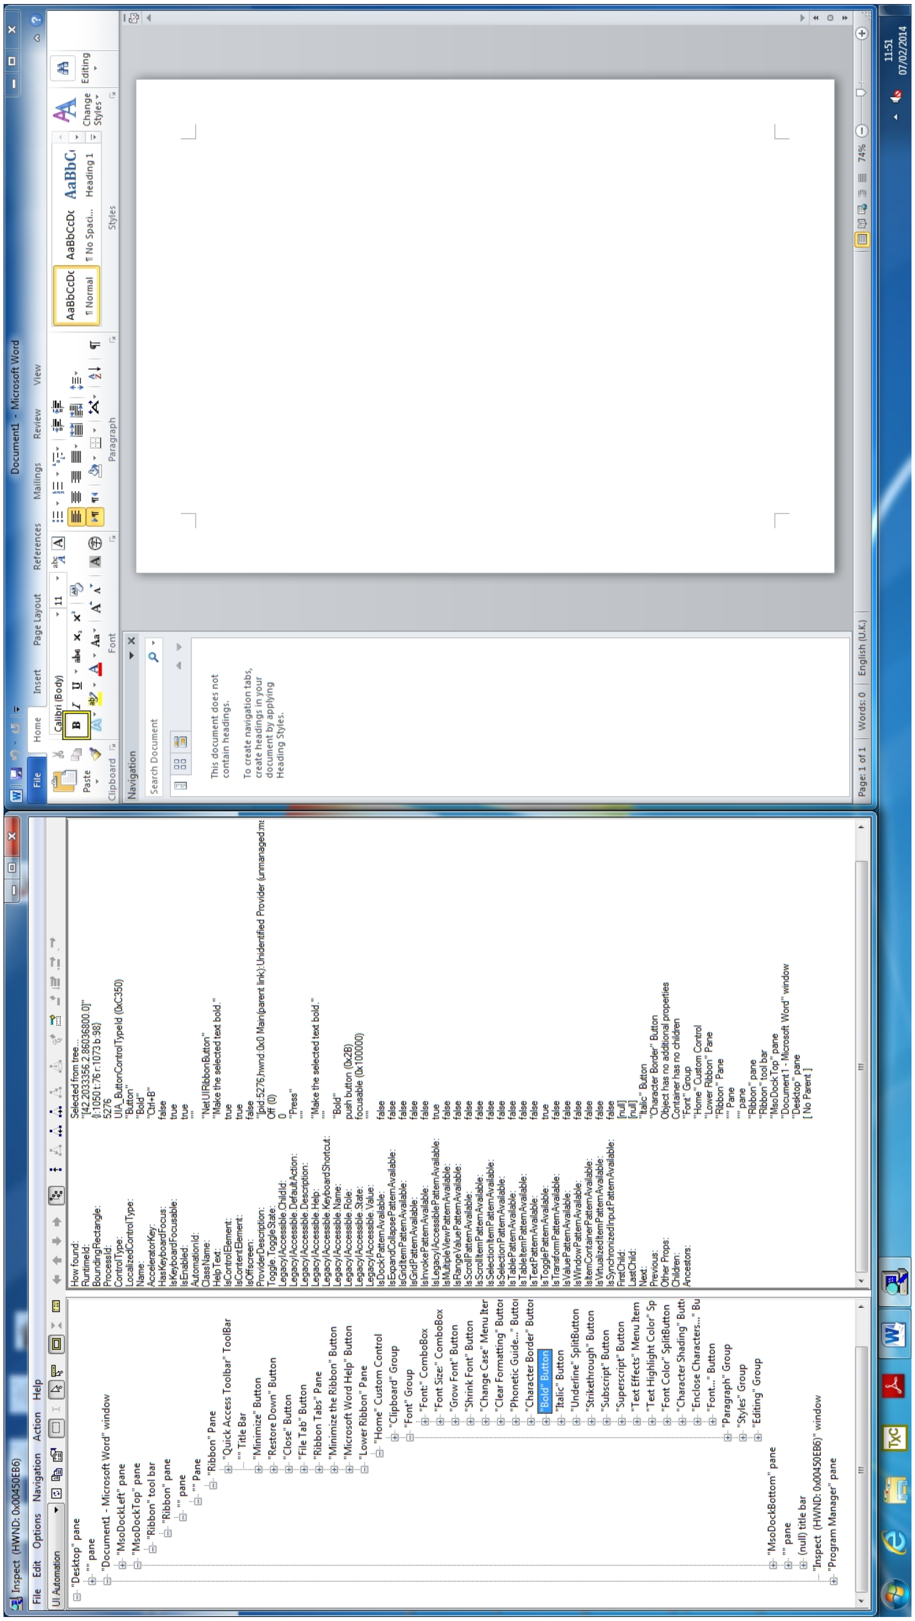

Figure B.2: Inspect (Inspect.exe) is a Windows-based tool that enables the user to select any UI element and view the accessibility data. The UI Automation tree is visible for the highlighted “Bold” button within a Microsoft Word application.

There are multiple ways for clients to programmatically obtain UI Automation elements within the provider's application. For example, using the `FindAll` method with a condition statement, using the `TreeWalker` class to traverse the entire UI Automation tree or a subset of it, use the window handle (hWnd) of the control, its screen location, etc. According to the Microsoft documentation, the `TreeWalker` class tends to be faster for Win32 controls but the `FindAll` method is faster for WPF controls [UIA14]. The `FindAll` method has been employed given that WPF was released more recently and is likely to be supported for longer. Once the interface elements have been located they are stored in a list with their relevant properties such as the bounding rectangle, name and control type. UI Automation allows clients to subscribe to events within the provider's application. The element list is updated when the application window is moved or resized so that the haptic assistance can be repositioned appropriately. If the application window loses focus then the haptic cues are temporarily disabled until the focus is restored.

C# was chosen as the preferred language for developing the feature extractor due to the greater available literature. One of the major benefits of using the Microsoft .NET Framework is that it provides a language-independent development environment. Classes and libraries can be written in Visual Basic, C++, C#, etc. and used within the other .NET languages. A dynamic link library (.dll) has been created in C# that implements the feature extractor described in this section. The .dll is then imported into the C++ application and used to automate haptic assistance for existing software. Figure B.3 shows the successful generation of haptic assistance for Microsoft Word and Microsoft Paint.

B.3 Map generator

There may be instances where the feature extractor may not be useful because the interface may not contain UI Automation elements. The map generator presented in this section allows haptic maps to be manually created for applications that are not supported by the feature extractor. This will provide access to other software such as computer games, puzzles, etc. The maps are created by drag-and-dropping the mouse over features within the interface to indicate their size and location. The map generator has been designed predominantly for support workers because people with physical disabilities may find it difficult to accurately create a map without existing assistance. Once the map has been completed it is saved to file and can be reloaded each time that the interface is used. Figure B.4 shows a haptic map that has been created for a popular puzzle game used by NANSA participants called Thinkin' Things Collection 2.

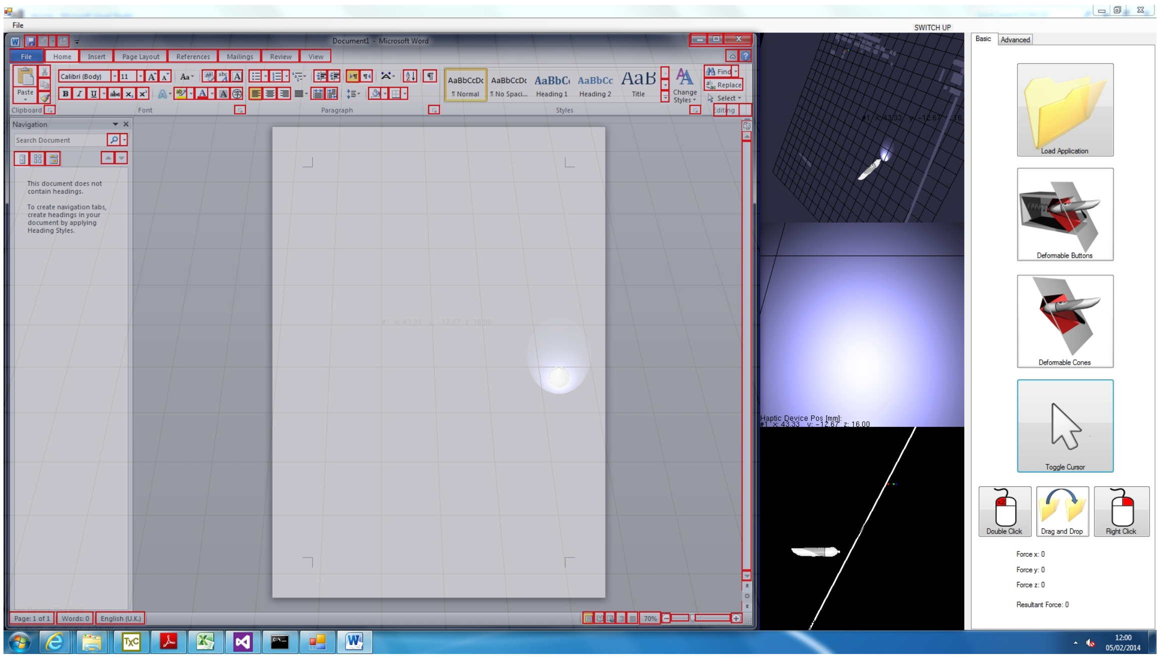

(a)

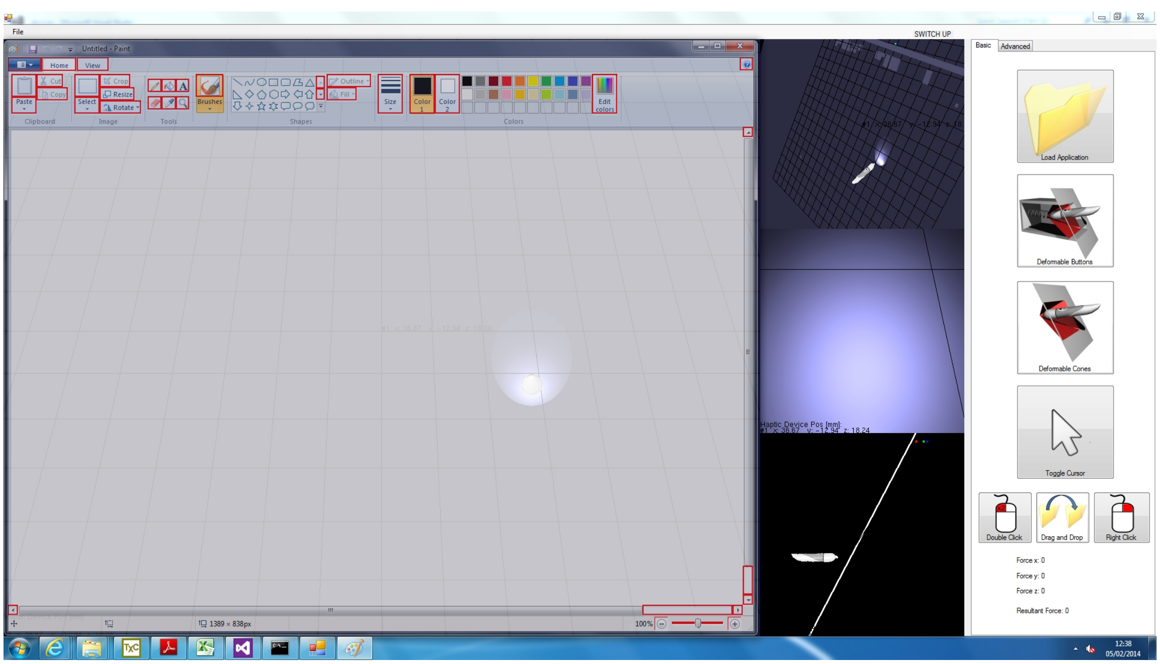

(b)

Figure B.3: The extracted features of a Microsoft Word interface automated with haptic assistance (a). The extracted features of a Microsoft Paint interface automated with haptic assistance (b).

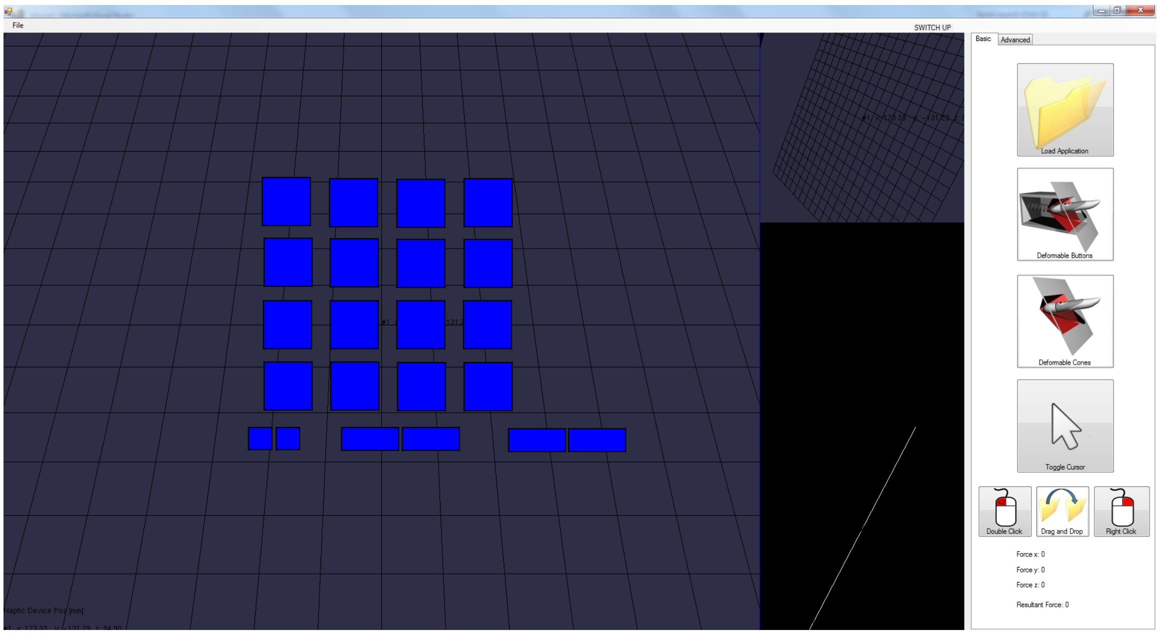

(a)

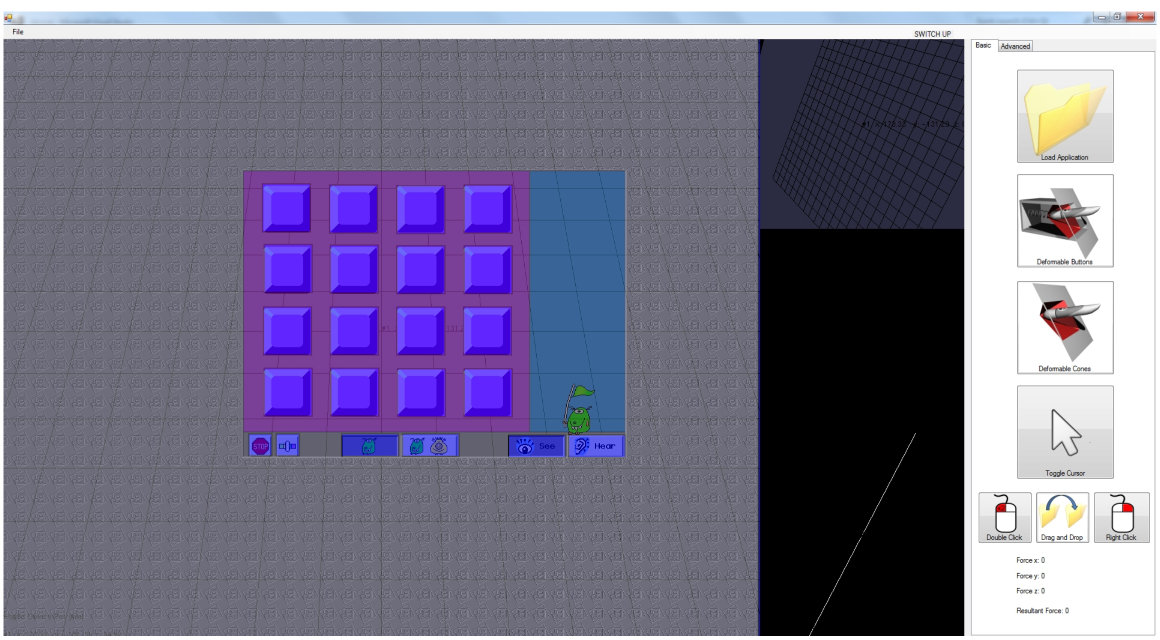

(b)

Figure B.4: A feature map manually created for the Thinkin' Things Collection 2 puzzle game (a). The exported feature map used to provide haptic assistance for the Thinkin' Things Collection 2 puzzle game (b).

Bibliography

[ADL11] C. T. Asque, A. M. Day, and S. D. Laycock. Haptic assisted target acquisition in a visual point-and-click task for computer users with motion-impairments. *IEEE Transactions on Haptics*, 5(2):120–130, 2011.

[ADL12] C. T. Asque, A. M. Day, and S. D. Laycock. Cursor navigation using haptics for motion-impaired computer users. In *Proceedings of the 2012 International Conference on Haptics: Perception, Devices, Mobility, and Communication - Volume Part I*, EuroHaptics'12, pages 13–24, 2012.

[AHL06] D. Ahlström, M. Hitz, and G. Leitner. An evaluation of sticky and force enhanced targets in multi target situations. In *Proceedings of the 4th Nordic Conference on Human-Computer Interaction: Changing Roles*, NordiCHI '06, pages 58–67, 2006.

[AMH95] M. Akamatsu, I. MacKenzie, and T. Hasbrouc. A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device. In *Ergonomics 38*, pages 816–827, 1995.

[ASK+05] T. Asano, E. Sharlin, Y. Kitamura, K. Takashima, and F. Kishino. Predictive interaction using the Delphian Desktop. In *Proceedings UIST 2005, ACM Press*, pages 133–141, 2005.

[AZ97] J. Accot and S. Zhai. Beyond Fitts' law: Models for trajectory-based HCI tasks. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI '97, pages 295–302, 1997.

[BGBL04] R. Blanch, Y. Guiard, and M. Beaudouin-Lafon. Semantic pointing: Improving target acquisition with control-display ratio adaptation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, pages 519–526, 2004.

[BGF02] M. Betke, J. Gips, and P. Fleming. The camera mouse: Visual tracking of body features to provide computer access for people with severe disabilities. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 10:1–10, 2002.

[BI03] R. Bates and H. Istance. Why are eye mice unpopular? A detailed comparison of head and eye controlled assistive technology pointing devices. *Universal Access in the Information Society*, pages 280–290, 2003.

[Bie90] E. A. Bier. Snap-dragging in three dimensions. In *Proceedings of the 1990 Symposium on Interactive 3D graphics*, I3D '90, pages 193–204, 1990.

[BJ95] E. Bergman and E. Johnson. Toward accessible human-computer interaction. In *Advances in Human-Computer Interaction (vol. 5)*, pages 87–113, 1995.

[BPK05] B. Bayart, A. Pocheville, and A. Kheddar. An adaptive haptic guidance software module for I-TOUCH: Example through a handwriting teaching simulation and a 3D maze. In *Proceedings of the Haptic Audio Visual Environments and their Applications.*, page 6, Ottawa, Canada, October 1-2 2005.

[Bre98] S. Brewster. The design of sonically enhanced widgets. *Interacting with Computers*, 11(2):211–235, 1998.

[BST88] B. Bied Sperling and T. S. Tullis. Are you a better "mouser" or "trackballer"?: A comparison of cursor-positioning performance. *SIGCHI Bulletin*, 19(3):77–81, 1988.

[BSZY11] S. Banihashem, S. Shishehchi, N. Zin, and N. Yatim. Accessible targets for motion-impaired users with hidden click zone technique. In *2011 International Conference on Pattern Analysis and Intelligent Robotics (ICPAIR)*, volume 2, pages 188–191, 2011.

[Bus96] G. Busby. Technology for the disabled and why it matters to you. In *Proceedings of the 5th International Conference on Computers Helping People with Special Needs. Part I*, ICCHP '96, pages 99–105, 1996.

[Cas92] S. P. Casali. Cursor control device use by persons with physical disabilities: Implications for hardware and software design. In *Proceedings of the Human Factors and Ergonomics Society*, Proceedings HFES'92, pages 311–315, 1992.

[CB05] A. Cockburn and S. Brewster. Multimodal feedback for the acquisition of small targets. *Ergonomics*, 48:1129–1150, 2005.

[CBM+12] F. Conti, F. Barbagli, D. Morris, C. Sewell, and S. Grange. CHAI3D open-source haptic API. `http://www.chai3d.org/`, 2012.

[CEB87] S. K. Card, W. K. English, and B. J. Burr. Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys, for text selection on a CRT. In *HCI*, pages 386–392, 1987.

[CLA14] Sensable Technologies - Claytools. `http://www.sensable.com/products-claytools-system.htm`, 2014.

[CVPC07] G. Casiez, D. Vogel, Q. Pan, and C. Chaillou. Rubberedge: Reducing clutching by combining position and rate control with elastic feedback. In *Proceedings of UIST*, pages 129–138, 2007.

[CZMM99] C. S. Campbell, S. Zhai, K. W. May, and P. P. Maglio. What you feel must be what you see: Adding tactile feedback to the trackpoint. *Proceedings of IFIP Conference on Human-Computer Interaction*, pages 383–390, 1999.

[Daw06] M. Dawe. Desperately seeking simplicity: How young adults with cognitive disabilities and their families adopt assistive technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 1143–1152, 2006.

[DJ06] J. T. Dennerlein and P. Johnson. Changes in upper extremity biomechanics across different mouse positions in a computer workstation. In *Ergonomics*, pages 1456–1469, 2006.

[DKM99] S. A. Douglas, A. E. Kirkpatrick, and I. S. MacKenzie. Testing pointing device performance and user assessment with the ISO 9241, Part 9 Standard. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '99, pages 215–222, 1999.

[DLB+05] L. Dominjon, A. Lécuyer, J. Burkhardt, G. Andrade-Barroso, and S. Richir. The "Bubble" technique: Interacting with large virtual environments using haptic devices with limited workspace. In *World Haptics Conference (joint Eurohaptics Conference and Haptics Symposium)*, pages 639–640, 2005.

[DM94] S. A. Douglas and A. K. Mithal. The effect of reducing homing time on the speed of a finger-controlled isometric pointing device. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '94, pages 411–416, 1994.

[DNS14] Dragon NaturallySpeaking. `http://www.nuance.com/dragon/index.htm`, 2014.

[DW02] C. Doerrer and R. Werthschuetzky. Simulating push-buttons using a haptic display: Requirements on force resolution and force-displacement curve. In *Proceedings of EuroHaptics*, 2002.

[DY01] J. T. Dennerlein and M. C. Yang. Haptic force-feedback devices for the office computer: Performance and musculoskeletal loading issues. *Human Factors*, 43(2):278–286, 2001.

[Epp86] B. W. Epps. Fitts' law models. *The Human Factors Society 30th Annual Meeting*, 30:327–331, 1986.

[EPP14] Microsoft - Enhanced Pointer Precision. `http://msdn.microsoft.com/en-us/library/windows/hardware/gg463319.aspx`, 2014.

[FF01] C. Fuhrer and S. Fridie. There's a mouse out there for everyone. In *Proceedings CSUN 01.*, 2001.

[FJS+10] L. Findlater, A. Jansen, K. Shinohara, M. Dixon, P. Kamb, J. Rakita, and J. O. Wobbrock. Enhanced area cursors: Reducing fine pointing demands for people with motor impairments. In *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*, UIST '10, pages 153–162, 2010.

[GMD09] C. Gunn, W. Muller, and A. Datta. Performance improvement with haptic assistance: A quantitative assessment. In *World Haptics 2009*, pages 511–516, Washington DC, USA, 2009.

[GMT14] The Phantom Omni haptic device. `http://www.geomagic.com/en/products/phantom-omni/overview`, 2014.

[GWW08] K. Gajos, J. Wobbrock, and D. Weld. Improving the performance of motor-impaired users with automatically-generated, ability-based interfaces. In *Proceedings of SIGCHI Conference on Human Factors in Computing Systems*, pages 1257–1266, New York, NY, USA, April 5-10 2008.

[HH08] B. Holbert and M. Huber. Design and evaluation of haptic effects for use in a computer desktop for the physically disabled. In *Proceedings of the 1st International Conference on Pervasive Technologies Related to Assistive Environments*, pages 9:1–9:8, Athens, Greece, 2008.

[HKL+01] F. Hwang, S. Keates, P. Langdon, J. Clarkson, and P. Robinson. Perception and haptics: Towards more accessible computers for motion-impaired users. In *Proceedings of the 2001 Workshop on Perceptive User Interfaces*, PUI '01, pages 1–9, 2001.

[HKLC03a] F. Hwang, S. Keates, P. Langdon, and J. Clarkson. Gravity well visibility in haptic interfaces for motion-impaired users. In *Proceedings of the 25th Annual International Conference of the IEEE EMBS*, pages 1670–1673, 2003.

[HKLC03b] F. Hwang, S. Keates, P. Langdon, and J. Clarkson. Multiple haptic targets for motion-impaired users. In *Proceedings of the Conference on Human Factors in Computing Systems*, pages 41–48, 2003.

[HLKC01] F. Hwang, P. Langdon, S. Keates, and J. Clarkson. Haptic assistance to improve computer access for motion-impaired users. In *Eurohaptics 2001 Conference Proceedings*, pages 176–178, 2001.

[HLKC03] F. Hwang, P. Langdon, S. Keates, and J. Clarkson. The effect of multiple haptic distractors on the performance of motion-impaired users. In *6th ERCIM Workshop*, pages 14–25, Italy, 2003.

[HLM+06] S. Harada, J. A. Landay, J. Malkin, X. Li, and J. A. Bilmes. The Vocal Joystick: Evaluation of voice-based cursor control techniques. In *Proceedings of the 8th International ACM SIGACCESS Conference on Computers and Accessibility*, Assets '06, pages 197–204, 2006.

[HMDH07] A. Hurst, J. Mankoff, A. K. Dey, and S. E. Hudson. Dirty desktops: Using a patina of magnetic mouse dust to make common interactor targets easier to select. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology*, UIST '07, pages 183–186, 2007.

[HNPD10] J. Hourcade, C. M. Nguyen, K. B. Perry, and N. L. Denburg. Pointassist for older adults: Analyzing sub-movement characteristics to aid in pointing tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1115–1124, 2010.

[Hwa03] F. Hwang. Partitioning cursor movements in point-and-click tasks. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '03, pages 682–683, 2003.

[Jac91] R. J. K. Jacob. The use of eye movements in human-computer interaction techniques: What you look at is what you get. *ACM Transactions on Information Systems*, pages 152–169, 1991.

[KD02] A. Kintsch and R. Depaula. A framework for the adoption of assistive technology. In *ASSETS 2002*, pages 1–10, 2002.

[KHL+02] S. Keates, F. Hwang, P. Langdon, P. Clarkson, and P. Robinson. The use of cursor measures for motion-impaired computer users. *Universal Access in the Information Society*, 2(1):18–29, 2002.

[KI03] M. Kobayashi and T. Igarashi. Considering the direction of cursor movement for efficient traversal of cascading menus. In *UIST*, pages 91–94, 2003.

[KLA+03] H. Koester, E. LoPresti, G. Ashlock, W. McMillan, P. Moore, and R. Simpson. Compass: Software for computer skills assessment. In *CSUN's 18h Annual Conference "Technology and Persons with Disabilities"*, Los Angeles, CA, 2003.

[KLCR00] S. Keates, P. Langdon, J. Clarkson, and P. Robinson. Investigating the use of force feedback for motion-impaired users. In *6th ERCIM Workshop*, pages 207–212, Italy, 2000.

[KLS05] H. Koester, E. LoPresti, and R. Simpson. Toward goldilocks' pointing device: Determining a "just right" gain setting for users with physical impairments. In *Proceedings of ACM SIGACCESS*, pages 84–89, Baltimore, MD, USA, 2005.

[KLS06] H. Koester, E. LoPresti, and R. Simpson. Factors influencing user performance with pointing devices. In *Proceedings of RESNA 29th International Annual Conference*, 2006.

[Koe03] H. Koester. Abandonment of speech recognition by new users. In *Proceedings of the 26th Annual Conference of the Rehabilitation Engineering and Assistive Technology Society of North America (RESNAU03).*, volume 3, 2003.

[KPW07] M. Kumar, A. Paepcke, and T. Winograd. Eyepoint: Practical pointing and selection using gaze and keyboard. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 421–430, 2007.

[KS02] A. S. Karimullah and A. Sears. Speech-based cursor control. In *Proceedings of the Fifth International ACM Conference on Assistive Technologies*, pages 178–185, 2002.

[KT05] S. Keates and S. Trewin. Effect of age and Parkinson's disease on cursor positioning using a mouse. In *Proceedings of the 7th International ACM SIGACCESS Conference on Computers and Accessibility*, Assets '05, pages 68–75, New York, NY, USA, 2005.

[KYM07] R. Kuber, W. Yu, and G. McAllister. Towards developing assistive haptic feedback for visually impaired internet users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 1525–1534, 2007.

[Lan00] C. Lankford. Effective eye-gaze input into windows. In *Proceedings of the 2000 Symposium on Eye Tracking Research & Applications*, ETRA '00, pages 23–27, 2000.

[LBA00] E. LoPresti, D. Brienza, and J. Angelo. Computer head control software to compensate for neck movement limitations. In *Proceedings Conference on Universal Usability*, pages 147–148, Arlington, Virginia, United States, 2000.

[LBAG03] E. F. LoPresti, D. M. Brienza, J. Angelo, and L. Gilbertson. Neck range of motion and use of computer head controls. *Journal of Rehabilitation Research and Development*, pages 199–212, 2003.

[LCR07] E. Lank, Y. Cheng, and J. Ruiz. Endpoint prediction using motion kinematics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 637–646, New York, NY, USA, 2007.

[LD07] S. Laycock and A. Day. A survey of haptic rendering techniques. *Computer Graphics Forum*, 26:50–65, 2007.

[LHK+02a] P. Langdon, F. Hwang, S. Keates, J. Clarkson, and P. Robinson. Developing assistive interfaces for motion-impaired users using cursor movement analysis in conjunction with haptic feedback. In *Proceedings ICDVRAT '02.*, pages 223–230, Reading, UK, 2002.

[LHK+02b] P. Langdon, F. Hwang, S. Keates, P. Clarkson, and P. Robinson. Investigating haptic assistive interfaces for motion-impaired users: Force-channels and competitive attractive-basins. In *Proceedings of Eurohaptics*, pages 122–127, UK, 2002.

[LHPO09] Y. Li, J. Huegel, V. Patoglu, and M. O'Malley. Progressive shared control for training in virtual environments. In *Proceedings of the World Haptics 2009 - Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 332–337, 2009.

[LKCR00] P. Langdon, S. Keates, J. Clarkson, and P. Robinson. Using haptic feedback to enhance computer interaction for motion-impaired users. In

Proceedings of ICDVRAT'00, pages 25–32, Alghero, Italy, September 2000.

[LKS02] M. D. Laehyun Kim, Anna Kyrikou and G. S. Sukhatme. An implicit-based haptic rendering technique. In *International Conference on Intelligent Robots and Systems*, pages 2943–2948, EPFL, Switzerland, October 2002.

[LS80] D. Lee and B. Schachter. Two algorithms for constructing a Delaunay triangulation. In *Proceedings of the International Journal of Parallel Programming*, pages 219–242, 1980.

[LS05] J. Levine and M. Schappert. A mouse adapter for people with hand tremor. *IBM Systems Journal*, 44(3):621–628, August 2005.

[MB05] M. J. McGuffin and R. Balakrishnan. Fitts' law and expanding targets: Experimental studies and designs for user interfaces. *ACM Transactions on Computer-Human Interaction*, pages 388–422, December 2005.

[MBG+08] J. J. Magee, M. Betke, J. Gips, M. R. Scott, and B. N. Waber. A human-computer interface using symmetry between eyes to detect gaze direction. *IEEE Transactions on Systems, Man, and Cybernetics. Part A*, 38(6):1248–1261, 2008.

[MC95] P. Millman and J. Colgate. Effects of non-uniform environment damping on haptic perception and performance of aimed movements. In *Proceedings of Eurohaptics Conference*, pages 703–712, 1995.

[MD96] A. K. Mithal and S. Douglas. Differences in movement microstructure of the mouse and the finger-controlled isometric joystick. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 300–307, 1996.

[MG07] C. Mauri and T. Granollers. On the assessment of the interaction quality of users with cerebral palsy. In *Proceedings of the Second International Conference on Availability, Reliability and Security*, pages 799–805, 2007.

[MG08] R. Mandryk and C. Gutwin. Perceptibility and utility of sticky targets. In *Proceedings of Graphics Interface 2008*, pages 65–72, 2008.

[MGB01] M. R. McGee, P. Gray, and S. Brewster. Haptic perception of virtual roughness. In *CHI'01 Extended Abstracts on Human Factors in Computing Systems*, pages 155–156, 2001.

[MKS01] I. S. MacKenzie, T. Kauppinen, and M. Silfverberg. Accuracy measures for evaluating computer pointing devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '01, pages 9–16, 2001.

[MM99] B. Martin and L. McCormack. Issues surrounding assistive technology; use and abandonment in an emerging technological culture. In *Proceedings of Association for the Advancement of Assistive Technology in Europe (AAATE) Conference*, pages 413–20, 1999.

[MR02] P. Majaranta and K. Räihä. Twenty years of eye typing: Systems and design issues. In *Proceedings of the 2002 Symposium on Eye Tracking Research & Applications*, pages 15–22, 2002.

[MRF+96] W. Mark, S. Randolph, M. Finch, J. Van Verth, and R. Taylor, II. Adding force feedback to graphics systems: Issues and solutions. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, pages 447–452, 1996.

[MSB91] I. S. MacKenzie, A. Sellen, and W. A. S. Buxton. A comparison of input devices in element pointing and dragging tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '91, pages 161–166, 1991.

[MSWB04] J. J. Magee, M. R. Scott, B. N. Waber, and M. Betke. Eyekeys: A real-time vision interface based on gaze detection from a low-grade video camera. In *Proceedings of the 2004 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'04) Volume 10*, CVPRW '04, page 159, 2004.

[Mur98] A. Murata. Improvement of pointing time by predicting targets with a PC mouse. In *Proceedings of the International Journal of Human-Computer Interaction*, pages 23–32, 1998.

[MZ98] T. Miller and R. Zeleznik. An insidious haptic invasion: Adding force feedback to the X desktop. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*, pages 59–64, San Francisco, California, United States, 1998.

[MZ99] T. Miller and R. Zeleznik. The design of 3D haptic widgets. In *Proceedings of the 1999 Symposium on Interactive 3D graphics*, I3D '99, pages 97–102, 1999.

[NHS14] Cerebral palsy. `http://www.nhs.uk/Conditions/Cerebral-palsy/Pages/Introduction.aspx`, 2014.

[OABG02] I. Oakley, A. Adams, S. Brewster, and P. Gray. Guidelines for the design of haptic widgets. In *16th British HCI Group Annual Conference*, pages 195–211, 2002.

[Oak99] I. Oakley. Comparing haptic effects in a GUI. In *Proceedings of First PHANToM Users Research Symposium, PURS99*, 1999.

[OBG01] I. Oakley, S. Brewster, and P. Gray. Solving multi-target haptic problems in menu interaction. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, pages 357–358, 2001.

[OHT14] SensAble Technologies OpenHaptics Toolkit version 3.0 programmers guide. `http://www.sensable.com/documents/documents/OpenHaptics\_ProgGuide.pdf`, 2014.

[OMBG00] I. Oakley, M. McGee, S. Brewster, and P. Gray. Putting the feel in 'look and feel'. In *Proceedings of the SIGCHI*, pages 415–422, The Hague, The Netherlands, 2000.

[PDE99] *ISO Ergonomic requirements for office work with visual display terminals (VDTs): Part 9 : Requirements for non-keyboard input devices, International Organisation for Standardisation.* ISO, Geneva, 1999.

[PZL04] X. Peng, W. Zhangand, and M. C. Leu. Interactive solid modeling in a virtual environment with haptic interface. In *Proceedings of Virtual and Augmented Reality Applications in Manufacturing.*, pages 43–61, 2004.

[Ram95] C. Ramstein. A multimodal user interface system with force feedback and physical models. In *Proceedings of IFIP INTERACT'95: Human-Computer Interaction*, pages 157–162, Norway, 1995.

[RKK97] D. C. Ruspini, K. Kolarov, and O. Khatib. The haptic display of complex graphical environments. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, pages 345–352, 1997.

[RMI06] M. E. Rodgers, R. L. Mandryk, and K. M. Inkpen. Smart sticky widgets: Pseudo-haptic enhancements for multi-monitor displays. In *Proceedings OF SMART GRAPHICS06*, pages 194–205, 2006.

[ROH14] Evonik Industries - Rohacell Foam. `http://www.rohacell.com/product/rohacell/en/Pages/default.aspx`, 2014.

[RRW00] M. Riemer-Reiss and R. Wacker. Factors associated with assistive technology discontinuance among individuals with disabilities. *Journal of Rehabilitation*, 66(3), 2000.

[RT96] C. Riviere and N. V. Thakor. Effects of age and disability on tracking tasks with a computer mouse: Accuracy and linearity. *Journal of Rehabilitation Research and Development*, 33(1):6–15, March 1996.

[SBM+95] K. Salisbury, D. Brock, T. Massie, N. Swarup, and C. Zilles. Haptic rendering: Programming touch interaction with virtual objects. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics*, I3D '95, pages 123–130, 1995.

[SCO14] Introduction to Scope. `http://www.scope.org.uk/home/scope.shtml`, 2014.

[Shi93] K. B. Shimoga. A survey of perceptual feedback issues in dexterous telemanipulation. II. Finger touch feedback. In *Proceedings of the 1993 IEEE Virtual Reality Annual International Symposium*, VRAIS '93, pages 271–279, 1993.

[SHL09] M. Stocks, S. Hayward, and S. Laycock. Interacting with the biomolecular solvent accessible surface via a haptic feedback device. In *BMC Structural Biology 9*, page 69, 2009.

[SSC99] M. Smith, J. Sharit, and S. Czaja. Aging, motor control, and the performance of computer mouse tasks. *Human Factors*, 41(3):389–396, 1999.

[TKM06] S. Trewin, S. Keates, and K. Moffatt. Developing steady clicks: A method of cursor assistance for people with motor impairments. In *Proceedings of ASSETS'06*, pages 26–33, Portland, Oregon, USA, 2006.

[TMI14] TalkItMouseIt 2. `http://www.amazon.com/TalkItMouseIt-2/dp/B00B1MXHF6`, 2014.

[TP99] S. Trewin and H. Pain. Keyboard and mouse errors due to motor disabilities. *International Journal of Human-Computer Studies*, 50(2):109–144, 1999.

[UIA14] Microsoft - Find a UI Automation Element for a List Item. `http://msdn.microsoft.com/en-us/library/ms788741(v=vs.110).aspx`, 2014.

[WFL09] J. Wobbrock, J. Fogarty, and S. Liu. The angle mouse: Target-agnostic dynamic gain adjustment based on angular deviation. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, pages 1401–1410, Boston, MA, USA, April 4-9 2009.

[WG07] J. Wobbrock and K. Gajos. A comparison of area pointing and goal crossing for people with and without motor impairments. In *Proceedings of ACM SIGACCESS*, pages 3–10, Tempe, Arizona, USA, 2007.

[WM06] J. Wobbrock and B. Myers. Trackball text entry for people with motor impairments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 479–488, 2006.

[WPC+04] D. Weir, M. Peshkin, E. Colgate, P. Buttolo, J. Rankin, and M. Johnston. The haptic profile: Capturing the feel of switches. In *Proceedings of the 12th International Conference on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, HAPTICS'04, pages 186–193, 2004.

[WPS+02] S. Wall, K. Paynter, A. Shillito, M. Wright, and S. Scali. The effect of haptic feedback and stereo graphics in a 3D target acquisition task. In *Proceedings of Eurohaptics*, pages 23–29, Edinburgh, UK, 2002.

[WWBH97] A. Worden, N. Walker, K. Bharat, and S. Hudson. Making computers easier for older adults to use: Area cursors and sticky icons. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI '97, pages 266–271, 1997.

[YM10] V. Young and A. Mihailidis. Difficulties in automatic speech recognition of dysarthric speakers and implications for speech-based applications used by the elderly: A literature review. *Assistive Technology*, pages 99–112, 2010.

[ZCBLG03] S. Zhai, S. Conversy, M. Beaudouin-Lafon, and Y. Guiard. Human online response to target expansion. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, pages 177–184, 2003.

[ZKE08] M. Zöllner, H. Koesling, and K. Essig. Gaze-contingent human-computer interaction in view of the midas-touch problem. *Perception*, 37(Supplement):11–11, 2008.

[ZMI99] S. Zhai, C. Morimoto, and S. Ihde. Manual and gaze input cascaded (MAGIC) pointing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '99, pages 246–253, 1999.

[ZS95] C. Zilles and J. Salisbury. A constraint based God-object method for haptic display. In *Proceedings of the IEEE Conference on Intelligent Robots and Systems*, pages 146–151, 1995.