Optimising visuo-locomotor interactions in a motion-capture virtual reality rehabilitation system

by

Mathew Alan Gilbert

A thesis submitted to the University of York in fulfilment of the
requirements for the degree of Doctor of Philosophy

Department of Electronics
University of York

January 2013

Abstract

This thesis presents the research-driven design and development of Stromohab: A motion-capture virtual-reality locomotion simulator for the research and rehabilitation of gait disorders following stroke. Software and hardware components are designed, developed and tested to facilitate and motivate patients in rehabilitative interactive avatar-based locomotor tasks. The system is then used to investigate systematically on healthy volunteers the known problem of distance underestimation in virtual environments by testing and analysing all combinations of cross-planar translation of leg movement to avatar actuated movement in a virtual environment. Specific performance deficits in the sagittal plane are confirmed and compared to those from coronal and transverse motion. Potential improvements of adding in isolation monocular cues for perspective, illumination, or size, and binocular cues from 3D stereo anaglyphs, are investigated, leading to a proposed movement model and scaling solution that both explains and resolves the observed deficit empirically in a practical locomotor task. Overall, the findings demonstrate the importance for the design and application of virtual environment interfaces of quantifying the underlying mechanisms in order to ensure accurate and controlled reproduction of a user's movement. These would be of particular significance in medical rehabilitation for neurological patients, for whom consideration of cognitive load and the potential for improper re-adaptation when returning to real world environments can be critical. It is envisaged that this study will be useful to technologists, clinicians and other professionals who apply the rapidly developing, increasingly accessible and beneficial motion capture and virtual reality technologies to medicine and related applications.

Contents

List of Figures

List of Tables

Acknowledgements

Foremost, I would like to express my gratitude to my advisor and mentor Dr. Adar Pelah, for his support of my Ph.D study and research. For his encouragement, enthusiasm, patience and motivation. His scholarly input, guidance and belief in my abilities have continually encouraged me to challenge myself and grow both academically and personally, and the breadth of his work has contributed to a unique learning experience extending far beyond the realm of the research described in this document. Without his help in securing funding support for my postgraduate research, it is unlikely that I would have been able to pursue my academic and career aspirations to the extent I have over the course of this project.

Funding for this research was also provided by a Holbeck Charitable Trust scholarship, and I am thankful to them for their support of this project.

I also thank my second advisor, Dr. Martin Robinson, for his advice and support in our committee meetings, and in his role as chair of the Physical Sciences Ethics Committee.

With regards to my development as an academic, I'd like to acknowledge Dr. Dave Pearce. I count myself lucky to have worked alongside such an inspirational educator, whose teaching ability and style I respect greatly.

It is unlikely I would have travelled so far down this academic road without the unwavering belief and support of my mum, Jane Gilbert. Your unconditional faith in me, all the financial support (one day, you will be driving that TT!), and continued interest in my contributions to the world have immeasurably contributed to the person I am today, and I will be forever thankful to you, the wonderful parent you are.

This thesis would also not have been possible without my wife, Becky Gilbert. Your understanding and support throughout the late nights, early mornings and working weekends is beyond comparison, and your boundless statistical knowledge, academic dexterity, intellectual curiosity, and above all, your love, contribute immeasurably to my life every single day. Your proficiency in ensuring the constant availability of doughnuts during the final months of this project was also nothing short of incredible, and I'm happy that somehow, we've both survived this journey with our sanity (relatively) intact.

While I cannot possibly mention everyone who has contributed to my development as an academic and a person, the enduring amusement of my sister Rachel has kept me grounded during the plentiful risible moments in my life, and the

transatlantic support of Deb and Gordon has kept a smile on my face during some of the tougher times. Their support and contributions to the organisation of the wedding of two graduate students is but one demonstration of the wonderful people they are.

Finally, this paper is dedicated to my Grandfather, whose ineffable personality handled the effects of Parkinson's and Stroke with a smile and a flick of his James Bond walking stick.

For you, Grandad, and the magic tree.

Declaration

I hereby certify that I am the sole author of this thesis and that no part of this thesis has been published or submitted for publication.

I certify that, to the best of my knowledge, this thesis does not infringe upon anyones copyright nor violate any proprietary rights and that any ideas, techniques, quotations, or any other material from the work of other people included in my thesis, published or otherwise, are fully acknowledged in accordance with the standard referencing practices.

Chapter 1

Introduction

This thesis contributes new knowledge relevant to the design of Virtual Reality (VR) environments that are used in conjunction with motion-capture based control interfaces. It presents an exploration of the current and potential applications of VR and locomotion simulators in post stroke rehabilitation, focusing specifically on lower-body movement and gait rehabilitation. It addresses the need for a detailed assessment of a technology that is rapidly gaining use within many rehabilitation fields, despite a lack of consistent and congruent methods and perspectives (Keshner, 2004; Crosbie et al., 2007; Braccialli et al., 2012). The research explores a perception that past and current VR based rehabilitation systems are not designed or developed with clinical use in mind, and hence present limited possibilities for widespread adoption and maximum useful rehabilitative impact (Schultheis et al., 2002; Edmans et al., 2006; **?**). This exploration requires a detailed review of the current state of the art, and also leads to the research-driven design and implementation of a gait rehabilitation system for stroke survivors that presents solutions to some of the practical and technical issues surrounding the development of such systems.

The complete design and implementation of such a system naturally raises many interesting research questions. After the complete implementation of an initial research platform, I focus on the unsolved problem of distance estimation in virtual environments. A systematic investigation of incremental solutions to the problem is conducted, concluding with a solution that is simple to implement and shown to be effective in multiple virtual environment scenarios.

The research also examines the integration of the developed system into a clinical environment in readiness for clinical trials that are to be conducted in the near future. A case study reveals residual issues of technology usage and adoption, and in order to further lower the barriers to clinical adoption of this useful technology, suggestions applicable to this and other VR based rehabilitation systems are proposed.

An additional problem in utilising VR environments in rehabilitation is that of patient engagement. In the wider field of VR, one driver of engagement is a sense of 'presence' in the environment. It is posited that using familiar, realistic environments may increase a patient's engagement with a virtual environment, and this leads to an investigation into the feasibility of using high definition video footage of suitable locations as a virtual environment. This approach is more analogous with the field of Augmented Reality (AR) than VR, and thus background and a small case study revealing the technological challenges associated with this approach are presented.

As this thesis is intended to serve as a reference for both engineering and therapeutic collaborators the language within is made to be accessible to a number of specialities. Where technical or clinical terms, or common abbreviations are used, brief definitions may be found in footnotes, or the glossary.

For the readers' reference, Figure 1.1 illustrates an early Stromohab system design.

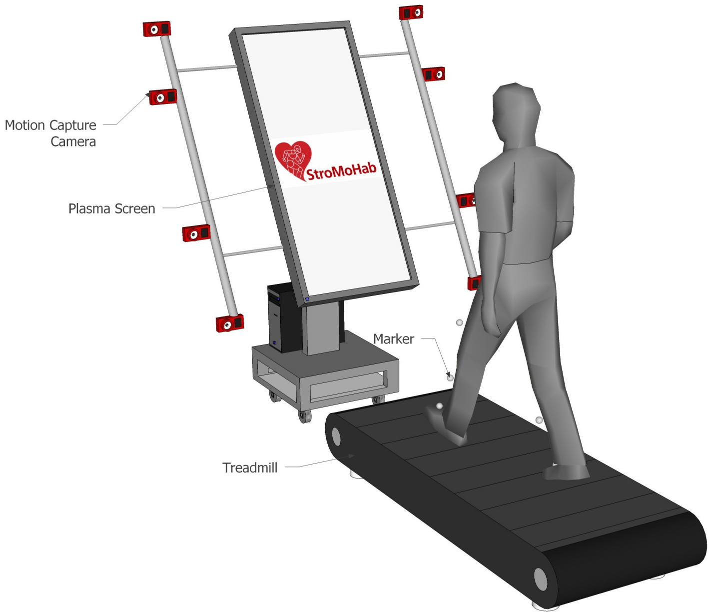

Figure 1.1: *An early Stromohab concept design.*

1.1 Motivation and Scope

The scope of this research spans a number of fields. I present a critical review of the current state of research pertaining to the use of both motion capture and VR in stroke rehabilitation. Recent advances in motion-controlled human-computer interfaces are discussed, and I then further focus on the use of these technologies in lower-body and gait rehabilitation, showing some of the issues slowing the wider uptake of these technologies in therapeutic settings. A brief overview of the implemented research platform is described from a high-level engineering perspective, before the presentation of a series of psychophysical studies, which both validate the platform as a useful, effective research tool, and quantify and resolve a common source of motion accuracy errors made within virtual environments.

The motivation for this research stems from a number of sources. Initial laboratory projects (Dring, 2007; Holmes, 2007; Grimes, 2008) had explored the development of an integrated treadmill-based rehabilitation platform, and established a laboratory based prototype system. While this had demonstrated the potential of an improved VR enhanced treadmill-based locomotion simulator over previous technologies including Weiss et al. (2004a); Crosbie et al. (2007); Pelah et al. (2009), the system required extensive adaptations for use in a clinical environment. Running on a platform consisting of multiple PCs, using webcams attached to a bulky scaffold surrounding the treadmill, preparing the system for use was both complex and time-consuming. In short, the prototype system lacked many desirable characteristics as described in Chapter 2, but had effectively demonstrated the potential of the paradigm.

More generally, an initial survey of the literature shows that, while a plethora of upper-body rehabilitation devices utilising virtual reality have been developed, there is a lack of lower-body devices using the same technologies. Additionally, the lower-body focused devices that have been developed tend towards a high financial cost You et al. (2005); Mirelman et al. (2009), or require long and/or complex setup procedures (Riener and Lünenburger, 2010), raising the barriers to adoption and use by therapeutic staff , Schultheis and Rizzo (2001); Rose et al. (2005), and Crosbie et al. (2007) for review).

Finally, the issue of distance underestimation is explored using the implemented system. As far as is known, though virtual environment literature often notes a similar observation (Rieser et al., 1990; Witmer and Kline, 1998; Servos, 2000; Knapp and Loomis, 2004), few in-depth attempts to quantify the phenomenon have been made (Richardson and Waller, 2005; Mohler et al., 2006). With the use of both motion-capture and virtual environment technologies proliferating as factors such as price and required technical expertise approach consumer levels, the quantification and resolution of this issue becomes increasingly important. In cases where these technologies are utilised in a rehabilitative environment, such as Jack et al. (2001); Merians et al. (2002); Boian et al. (2002); Broeren et al. (2004), within which the significance of properties such as patient engagement, cognitive load and ecological validity is greatly increased, a transparent, accurate, coherent and demonstrable resolution to this problem is not only important, but arguably essential.

1.2 Stroke

An established and universally accepted definition of stroke is 'acute neurological dysfunction of vascular origin... with symptoms and signs corresponding to the involvement of focal areas of the brain' (World Health Organisation, 1989). It can be described as the rapid onset of neurological deficits caused by a disruption in the blood supply to the brain. This disruption can be caused by either an obstruction in the blood flow (Ischemic stroke), or the rupture of an artery that supplies the brain (Hemorrhagic stroke). Ischemic stroke occurs either when a thrombus, or blood clot, forms and obstructs blood flow to a part of the brain, or when a blood clot forms elsewhere in the body and breaks off into the bloodstream (an embolus), before being carried to the brain and obstructing the blood supply. Hemorrhagic stroke occurs when an artery in the brain bursts (cerebral hemorrhage), or when a blood vessel on the brain's surface ruptures (subarachnoid hemorrhage). Ischemic stroke accounts for approximately three-quarters of all strokes (Lloyd-Jones et al., 2009). Either form of blood flow disruption leads to lack of blood flow to a part of the brain, killing brain cells due to a lack of oxygen and leading to symptoms that reflect the specific or acute damage to the brain. Brainin and Heiss (2009) lists the key symptoms of acute stroke as including:

- Dizziness, trouble walking, loss of balance and perception.
- Speech difficulties or deficits.
- Numbness, weakness or paralysis on one side of the body (hemiparesis).
- Blurred, darkened, or double vision.
- Sudden severe headache.

A stroke is distinguished from a transient ischemic attack (TIA) by the fact that neurological deficits caused by stroke persist longer than 24 hours (Easton et al., 2009).

Once an acute stroke victim has stabilised, symptoms that persist will vary from none, to many, dependant on the area of the brain and the amount of tissue that was damaged. As can be see in Figure 1.2, the brain consists of a number of specialised regions (Fodor, 1983; Caramazza and Coltheart, 2006). Each of these regions tends to have a different, specialised function, and as stroke can affect any area of the brain, a stroke victim may be left with major sensory, motor and/or cognitive problems.

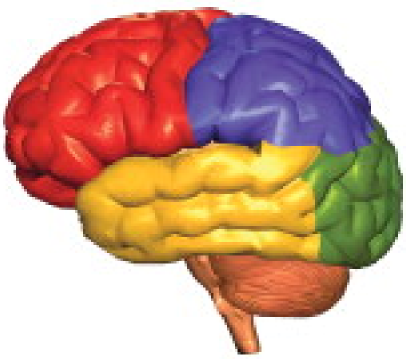

Figure 1.2: *Regions of the brain*[1]

As each hemisphere of the brain controls the opposite side of the body, a stroke affecting one side of the brain will generally result in neurological complications on the opposing side of the body. From Smeltzer et al. (2010), a stroke in the right side of the brain may produce:

- Paralysis on the left side of the body
- Vision field deficits
- Quick, inquisitive behavioural style
- Memory loss

A stroke in the left side of the brain may produce:

- Paralysis on the right side of the body
- Speech and/or language problems
- Slow, cautious behavioural style
- Memory loss

---

[1] Image from `http://www.strokeassociation.org`

There are many causes of stroke, though the most common risk factors include: high blood pressure or cholesterol, smoking, diabetes, atherosclerosis (thickening of the arteries) and obesity (Warlow et al., 2011). Party (2008) reports that approximately 15 percent of stroke victims die shortly after a stroke, and 10 percent of stroke survivors recover almost immediately. They additionally state that of the other 75 percent, 10 percent require care in a nursing home or other long-term care facility, 40 percent experience moderate to severe impairments requiring specialist care and 25 percent recover with minor impairments. Most of the three quarters will require movement rehabilitation of some kind.

1.2.1 Stroke Rehabilitation

Treatment after the acute stage of stroke focuses on improving function to maximal potential so that the stroke survivor can become as independent as possible. Within the brain, there are on average 86 billion neurons (Herculano-Houzel and Lent, 2005), areas of which die during stroke. The brain however, has the capacity to reorganise the pathways between neurons - indeed, Gopnik et al. (2000) describe neurons as growing telephone wires that communicate with one another. This capacity is called *neuroplasticity*, and on a neural level, stroke rehabilitation seeks to encourage the growth of new neural connections, as well as strengthen and adapt surviving neuronal pathways Bracewell (2003); Cramer and Riley (2008).

Multiple specialised clinical teams are usually involved in stroke rehabilitation, due to the wide range of deficits a stroke can cause. As every stroke is different, after the acute phase, there is no single or main treatment. Once a person is medically stable, rehabilitation may begin immediately. The sooner rehabilitation is started, the more likely a positive outcome - rehabilitation may begin as soon as 24 hours a stroke. Depending on the severity of the stroke, current rehabilitation options include:

- A hospital rehabilitation unit.
- A subacute care unit.
- Home therapy.
- Home with outpatient therapy.
- A long-term care facility providing therapy and nursing care.

Clinical treatments involved in stroke rehabilitation include physiotherapy, which aims to assess and treat problems with movement and balance, occupational therapy, which aims to help the stroke survivor regain competence in the activities of daily living, and speech therapy, which aims to manage swallowing and

/ or communication difficulties. While a stroke survivor may not require any of these treatments, as stated above, it is likely that most will require movement rehabilitation. This is generally delivered in the form of physiotherapy.

Physiotherapy

There are a number of different frameworks physiotherapists work within today, most of which have been developed since the 1950's (Jensen et al., 2000). Before then, neurological physiotherapy did not exist as a specialism, and treatment focused on functional performance, using the unaffected side of the body to compensate for any lack of abilities on the affected side. Though evidential data is lacking, this approach is now thought to have led to increased dysfunction on the affected side. Since the 1950's, new approaches built on neurophysiological theories have been developed (Bobath, 1990; Brunnstrom, 1956; Rood, 1954), and more recent work has introduced concepts from disciplines such as biomechanics, psychology and education (Affolter, 1981; Carr et al., 1982). While individual therapists may use a single, or a combination of these frameworks, the current most widely used framework is likely to be Bobath (Pollock et al., 2008), which is based on neurophysiological principles. Frameworks popular in specific countries include Carr and Shepherd, more prevalent in Australia, Pető (Hungary), and Affolter (Switzerland).

The Bobath approach states that postural control is the basis on which patients begin to develop their skills. Patients learn postures and movements; therapists analyse these and look for abnormalities. In treating a patient, the therapist promotes motor learning through the use of sensory information such as tactile cues via manual contact, and verbal directions. This sensory information is used to reinforce weak movement patterns and discourage overactive movement patterns. A full introduction to the Bobath concept and approach can be found in Bobath (1990).

Carr et al. (1982) developed a motor learning programme based on the retraining of functional motor control essential to everyday life. The main aim is to regain control of muscles, and focused, deliberate practice gives way to more automatic action as a patient improves. The treatment programme is broken down into a number of sections representing functions from everyday life, such as sitting up from supine, standing and sitting down, and walking.

Pető's method focuses on group rather than individual work, and chanting a 'rhythmical intention' such as 'I clasp my hands' is an integral part of the learning process (Cotton and Kinsman, 1983). Individual tasks are developed for a patients, through which they work to achieve a common group goal.

Affolter (1981) proposed the hypothesis that the adequacy of perceptual processes is a prerequisite for complex human performance, and the interaction between the person and the environment is viewed as fundamental to the learning process. A number of stages of recovery are suggested, characterised by levels of patient anticipation of functional steps to be taken. In order to elicit full recovery, patients need to experience success in following the correct sequence and achieving the goal.

An objective assessment of these, and other neurophysiological therapies clearly shows that there are many similarities between them, despite the apparent strong differentiation put forward by the theorists and their proponents. Studies comparing approaches have not yet provided any definite results. Stern and McDowell (1970) compared the techniques of Brunnström with 'conventional' physiotherapy, using two groups of stroke patients, and concluded that there was no significant difference in a number of outcome measures. Bobath and Rood approaches have also been compared (Lord and Hall, 1986), finding that while both improved function and motor performance in standardised tests, there were no significant differences between them. One of the most methodological comparative studies was completed by (Dickstein et al., 1986). Three groups were compared: 'conventional treatment', Brunnström, and Bobath, and it standardisation of treatment methods was attempt via a training period prior to the studies commencement. Again, there were no significant differences between the groups after a six week treatment period, according to a standardised outcome measure (Barthel Index). Earlier in the study, conventionally treated patients did start to walk earlier, though this early effect may be explained by the Bobath approach discouraging walking before normal balance and posture are achieved. At best, what can be concluded from these studies is that more treatment appears to provide better clinical outcomes than less, and that early, targeted treatment may provide a short-term advantage.

More recently, improved understanding of the biomechanical and neural mechanisms underpinning motor learning has lead to the development of evidence-based training and exercise protocols (Dean and Shepherd, 1997; Teixeira-Salmela et al., 1999), and there is a growing body of evidence showing that intensive, repeated training may be necessary to modify neural organisation (Kopp et al., 1999; Liepert et al., 1998) and effect recovery of functional motor skills (Taub et al., 1993). Further to this, it has been shown that three determinants of motor recovery are: early intervention, task-oriented training, and repetition intensity (Malouin et al., 2003) and Sveistrup (2004b) suggests that a major objective of rehabilitation is to identify the means to provide repeated opportunities for tasks that involve multimodal processes (sensory modalities including vision, audition, proprioception and haptics), to further enable increases in function.

1.3 Technology for Rehabilitation

Rehabilitation can be viewed as a dynamic process which uses available facilities to correct any undesired motion behaviour in order to reach an expectation such as ideal position (Sveistrup, 2004b). Zhou and Hu (2004) state that tracking human motion is vital in any rehabilitation scheme, postulating that in a rehabilitation course the movement of a stroke patient needs to be continuously monitored and rectified so as to hold a correct motion pattern (Zhou and Hu, 2008). This concept satisfies the requirement for a therapeutic environment to 'heighten sensory cues that inform the actor about the consequences of actions (forward modelling), and allow adaptive strategies to be sought (inverse modelling)' (Wann and Turnbull, 1993). However, avoiding overburdening a patient, as described by Moreland et al. (1998), by choosing what monitored variables to feedback and the sensory modalities through which to do so are important considerations in the design of any rehabilitation system, and therefore technologies used in a rehabilitation system must be carefully evaluated.

Biofeedback

Biofeedback can be defined as the use of of instrumentation to make covert physiological processes overt, and also includes electronic options for shaping appropriate responses (Basmajian, 1989; Moreland and Thomson, 1994; Dursun et al., 2004). While the neurological mechanisms underlying biofeedback are unclear, Basmajian (1982) suggests two possibilities: either new neural pathways are developed, or an auxiliary feedback loop recruits existing cerebral and spinal pathways. The latter explanation has gained some support in literature, with Wolf (1983) theorising that visual and auditory feedback activates unused or underused synapses in executing motor commands, and as such, establishes new neural networks and helps patients perform without feedback.

Since the 1960s, physiological sources used in biofeedback have included EMG (Prevo et al., 1982; Wolf and Binder-MacLeod, 1983; Burnside et al., 1982), joint angle (Koheil et al., 1980; Dursun et al., 1996), pressure or ground reaction force (Winstein et al., 1989; Wu, 1997) and position (Montoya et al., 1994; Aruin et al., 2000). Regardless of the sources used, feedback is indicated through visual display, auditory pitch or volume, or mechanical stimulation (Metherall et al., 1996).

Treadmill Training for Gait Rehabilitation

Post hemiplegic stroke, it has been suggested that treadmill training with partial body weight support and manual assistance may be a useful technique to facilitate functional walking (Mayr et al., 2007b). Although randomised clinical trials show inconsistent benefits of the application of this technique (Moseley et al., 2003), studies of motor learning and task-oriented training suggest that this approach may improve walking outcomes as compared to conventional physiotherapy (Dietz and Harkema, 2004; van Peppen, 2008), at least in part by presenting an opportunity for a patient to relearn motor patterns through repetitive progressive practice of the spatiotemporal and kinetic parameters of walking, such as limb loading, velocity, cadence and stride length (Hesse et al., 1994).

In addition to treadmill training, technological approaches to rehabilitation include robotic, or 'enforced' rehabilitation systems (Iwata, 2000), and one of the first commercial systems of this type was the *Gait Trainer*, developed by Hesse and Uhlenbrock (2000). It consisted of two footplates positioned on two rocking bars, attached to a doubled crank and rocker gear system. This meant that the system could generate a 3D motion for each leg, though as the trajectory of the footpads was fixed by the crank and gear system, extensive setup was required for each individual patient before every use. Yano et al. (2003) developed a similar, next-generation system, marketed as the *Gait-Master 2*. They minimised setup overhead by controlling 2 DOF motion platforms via easily configurable software, which meant that the system could be reprogrammed with individual patient gait profiles much more rapidly than the hardware-based Gait Trainer. The authors demonstrate some success in rehabilitating two hemiplegic patients with this system, though as is too common in studies of this type, no comparison with conventional rehabilitative methods is made, and there is little reference to the mechanisms of recovery.

The most widely used robotic rehabilitative system currently in use, is likely the *Hocoma Lokomat*[2], which operates on similar principles to the above systems. While an in depth review of robotic gait rehabilitation is out of the remit of this thesis, it is worth nothing that there is a large body of literature that is mostly supportive of the efficacy of such devices. For reference, a representative sample of this literature includes: Riener and Lünenburger (2010); Hidler et al. (2009);

Koenig and Wellner (2008); Krewer et al. (2007); Mayr et al. (2007a); Lunenburger and Colombo (2005, 2004); Jezernik and Colombo (2003). A *Lokomat* rehabilitation system currently costs approximately GB£120,000.

1.3.1 Virtual Environments

Before describing virtual reality, within the context of this study, it is worth considering how we perceive the physical world.

Our sense of physical reality is a construction from the symbolic, geometric and dynamic information presented to our senses (Ellis, 1991). Our perceptual systems are capable of interpreting both incomplete and noisy information - for example, when seeing only part of a whole object, through the utilisation of *a priori* knowledge, we know that it exists in a complete form (Gregory, 1968). We appear to store this *a priori* knowledge as controlled internal models (Kalman, 1960; Kleinman et al., 1970), using error feedback to update erroneous or out of date aspects of this knowledge. Therefore, much of our perception of the physical world is not, as might be assumed, an immediate consequence of the information from our sensory inputs, but rather a combination of incoming information and our existing internal models. We see, feel and hear things happen, assimilate the results of this, and then expect similar outcomes when experiencing similar sensory inputs in a future situation, while giving heightened attention when such expectations are not met.

Virtual environments elicit similar neural processing, most often utilising the visual sensory pathway by means of a display, usually head-mounted or fixed. Virtual environments may broadly be separated into three levels of virtualisation - *Space*, *Image*, and *Environment* (Ellis, 1991). *Space* is the most abstract of these levels, and can be viewed as the perception of a three-dimensional space from a flat surface that presents the pictorial cues for space - perspective, scaling, occlusion and texture gradients, as in a painting, or drawing. *Image* is the next most abstract, in which depth is perceived using accommodative cues[3], vergence cues [4], and optionally stereoscopic disparity[5]. A complete *Environment* combines all of the preceding depth information, and also adds viewpoint driven motion parallax[6], depth-of-field focus changes[7], and implements a wide

---

[2] `http://www.hocoma.com/fileadmin/user/Dokumente/Lokomat/bro_L6_120416_en_A4.pdf`. Last accessed Jan 2013.

[3] Visual focus in the eyes.

[4] Angles of gaze formed by the eyes.

[5] Discrepancies between left and right retinal images during binocular fusion.

[6] Apparent difference in the direction of movement or speed produced when a subject moves relative to the environment.

[7] Apparent blurring of objects outside of the focal field.

field of view, with no prominent frame. This additional information stimulates physiological reflexes such as accommodative vergence[8], the optokinetic reflex[9], and the vestibular-ocular reflex[10] (Erkelens and Collewijn, 1985a,b). This additional information, if implemented well, can significantly improve the sense of presence a user of a virtual environment will experience (Airey et al., 1990; Welch et al., 1996; Ijsselsteijn et al., 1998; Freeman et al., 1999; Sanchez-Vives and Slater, 2005). Due to this increased physiological stimulation however, an interactive virtual environment also has the capacity to cause unwanted side effects. Delays or lags in the environment for example, interfere with visual-motor adaptation (Held and Durlach, 1991; Cunningham et al., 2001b), and can produce motion sickness in system users (Erkelens and Collewijn, 1985b,a).

A complete virtual environment can be viewed as having three components - *content*, *geometry*, and *dynamics* (Ellis, 1991). Content can be further subdivided into *actors* and *objects*, where the difference lies in the ability of an actor to initiate interactions with other objects. We can also view the *self* as a special case of an actor, as it provides a viewpoint from which the environment can be constructed. For example, the balls on a snooker table may be described as the content of the snooker table environment, with the cue ball and snooker player as the self. When utilising a third person view, the self is often represented using an *avatar*[11]. The avatar is a primary source of biofeedback (Section 1.3), and may serve to increase activation of mirror neurons[12] (Billard and Arbib, 2002), regardless of the visual similarities between the avatar and the subject (Bray and Konsynski, 2007).

Geometry describes the scene of action (Ellis, 1991), which consists of *dimensionality, metrics and extent.* Dimensionality refers to the number of independent descriptive terms required to specify the position of every element within an environment. Metrics are a rule system that is applied to all described positions within an environment, to establish both element ordering within the environment (i.e. which objects are 'in front' or 'behind' other objects), and to define the concept of straight lines within the environment. The environment extent refers to the range of all possible values the element positions may describe. Dynamics are the rules of interaction of an environment's contents. Dynamics can range from the very simple: 'object 1 changes colour when touched by object

---

[8] Rotation of the eyes as result of the eyes focusing.

[9] Eye movement in response to objects moving in peripheral vision while the head is stationary.

[10] Eye motion in a direction opposing that of a head movement, resulting in the preservation of a stable image on the retina.

[11] A graphical representation of the user.

[12] Specialised brain cells that respond when we perform an action, *and* when we observe another person perform that action.

2', to highly complex, such as advanced motion and force simulations.

Presence within a virtual environment

Minsky (1980) introduces the concept of telepresence to describe the feeling that a human operator may have while interacting with a machine via a teleoperator system. The operator 'sees' what a remote machine is 'seeing', and uses their own limbs to control the machine's effectors. In this situation, an operator may develop a sense of being in a remote place - a sense of 'presence'. This experience of presence was thought to be conducive to operators task performance in the virtual environment.

This concept of telepresence has also been applied to experiences within virtual environments. Here, a user is immersed within an environment created and displayed using computer hardware, and they may be able to interact with the environment, in multiple ways. A feeling of being 'present' may develop in the same way Minsky (1980) noted for physical teleoperator systems.

As touched upon in Section 1.3.1, there are many factors that can affect the quality of experience a user of an interactive virtual environment may have. These include the *field-of-view* - the horizontal and vertical angles through which the scene is viewed. Using flat display technology, the field of view may be approximately 60°diagonal, compared to more than 180°horizontal and 120°vertical in normal human vision. Though less of an issue with modern high definition displays, the resolution of the display technology also affects the level of environmental detail that can be shown to a user, and as such can affect the quality of experience. The *frame-rate* of the system is also a factor - while the display will always update at a constant rate (generally between 60-80Hz), more complex virtual environments may lead to the system hardware not being able to maintain this speed of image generation. This results in non-smooth motion on the screen, and a jarring experience for the user. Frame-rate also affects system *latency* - the time between a user initiating an event (such as a limb movement), and the time that the system processes this event and responds. The recently introduced (Jackson, 2012) *HFR*[13] cinema is one reflection of the gaining appreciation for the importance of smooth temporal processing within even non-interactive media.

Therefore, to maximise a user's sense of presence within a virtual environment we must do all we can to ensure that the *field of view*, *frame-rate*, *latency*, are

---

[13] High frame rate, see Perry and Zorpette (2013) for an in-depth review.

as optimal as possible. It is worth noting however, that perceptual, physiological and cognitive aspects also contribute significantly to a sense of presence (Schuemie et al., 2001; Riva et al., 2004; Väljamäe, 2005; Steuer, 2006; Riecke et al., 2006). Further to this, the Stromohab platform uses a patented (Pelah, 2007) display system, in which the virtual environment is displayed to the subject on a tilted display screen, mounted in portrait format. This configuration is somewhat unusual, but has been used in a number of VR locomotion investigations (Dong et al., 2008; Casey et al., 2010). It also aligns with measurements of the amplitude of eye saccades while walking (Bahill et al., 1975; Luo et al., 2008), and as such may be most appropriate for locomotion tasks, as compared with a landscape configuration more suited for scanning and stationary visual search tasks (Chan and So, 2008; Rayner, 2009).

Virtual Environments for Rehabilitation

The investigation and use of virtual technology for neurological and physical rehabilitation has been steadily increasing in laboratories and clinics since the early 1990's (see Greenleaf and Tovar (1994); Kuhlen and Dohle (1995); Rose et al. (1996); Tarr and Warren (2002); Holden (2005) for comprehensive reviews and examples), and a growing body of work suggests that rehabilitation aiming for improvement in functional activities benefits from task-oriented biofeedback therapy (Wann and Turnbull (1993); Bradley et al. (1998); Moreland et al. (1998); Nichols (1997)). As such, (Duncan et al., 1995) states effective motor training must incorporate movement components, while (Schmidt and Wrisberg, 2007) find that a rehabilitation environment must resemble the targeted task in the relevant functional context. A VR environment allows a user to interact with a multidimensional, multisensory computer generated environment, in real-time (as defined by Wilson et al. (1996)), and VR has been shown to be effective in rehabilitation, as compared to conventional rehabilitation (Sveistrup et al., 2003). In similar work by the same author (Sveistrup et al., 2004), VR is shown to provide a unique medium where therapy can be provided within a functional, purposeful and motivating context. Carrozzo and Lacquaniti (1998) describe VR as allowing 'for the creation of a synthetic environment with precise control over a large number of physical variables that influence behaviour while recording physiological and kinematic responses', which suits requirements of a rehabilitation environment as described above. Keshner (2004) further states the case for the use of VR in rehabilitation, detailing advantages of the approach including the opportunity for:

- Ecological validity,

- Stimulus control and consistency
- Real-time diagnostic and rehabilitative feedback
- A safe testing and training environment
- Graduated exposure to stimuli
- The ability to distract or augment the patients attention
- Motivation for patients

VR is therefore highly suited to the creation of a suitable environment. Appropriate interaction implementation is also a crucial component of a rehabilitation system. Current-day technologies allow for the rendering of such an environment on standard hardware, using open-source development platforms such as OpenGL. While some aspects of VR in rehabilitation are yet to be proven, work has shown that is an effective biofeedback modality suitable for use in a rehabilitation context.

1.3.2 Motion Capture

In the context of movement rehabilitation, it follows that human motion tracking provides an ideal interface with a VR environment. The relevance of motion tracking in rehabilitation has been described in literature (Weiss et al., 2004b; Attygalle et al., 2008; Zhang et al., 2011), a summary of which is presented here.

There are a wide variety of human motion tracking technologies available today; Zhou and Hu (2008) classify them into three broad areas:

- Non-visual tracking, such as inertial or magnetic based systems.
- Visual tracking - visual marker based, marker-free visual based, and combinatorial systems that attempt to take advantage of the benefits of both techniques.
- Robot-aided tracking.

Visual marker based systems can be further classified as either active or passive systems. An active system allows the identification of individual markers placed on human body parts, simplifying processing of the incoming data. Limitations of active marker systems include cost, and movement restrictions due to cabling powering the markers. Passive visual marker based systems are generally less expensive, and merely require small markers to be placed on the body parts to be tracked, making them less intimidating for subjects/patients in a research or rehabilitation environment, as well as vastly decreasing session setup time. Davis (1991) provides an extensive review of both active and passive marker systems for applications in medical science, and both system types have been used extensively in research and rehabilitation (Charlton et al., 2004; Esquenazi and Mayer, 2004; Delahunt et al., 2007).

Overall, each type of system has advantages and disadvantages, as shown in Table 1.1.

Systems	Accuracy	Compactness	Computation	Cost	Drawbacks
Inertial	High	High	Efficient	Low	Drifts
Magnetic	Medium	High	Efficient	Low	Ferromagnetic Materials
Ultrasound	Medium	Low	Efficient	Low	Occlusion
Glove	High	High	Efficient	Medium	Partial Posture
Marker	High	Low	Inefficient	Medium	Occlusion
Marker-free	High	High	Inefficient	Low	Occlusion
Combinatorial	High	Low	Inefficient	High	Multidisciplinary
Robot	High	Low	Inefficient	High	Limited Motion

Table 1.1: *Performance comparison of different motion tracking systems reproduced from Zhou and Hu (2008)*

Welch and Foxlin (2002) state that modality-specific, measurement-specific, and circumstance-specific limitations affect the use of particular sensors in different environments. The Stromohab system (Chapter 2), aiming for a minimally invasive, low-cost method of tracking human motion, uses a passive visual marker based system[14]. Along with other techniques, optical motion capture systems suffer from sensor noise and require careful calibration (Delaney, 1998). Optical marker systems present a number of additional interesting technical challenges in data collection and movement reconstruction, and reconstructing human motion from motion capture data requires the implementation of various processing techniques, described below.

Human Motion Reconstruction

As described by Zordan and Van Der Horst (2003), during data collection and processing, detailed reconstruction of a subject's movements requires a nontrival mapping of the collected marker movement data, moving in Cartesian 3D-space, to a relative motion representation defined by joint angles plus a body centre or root. This problem (of determining a systems's kinematic parameters from the motion of the system) has been studied in depth in the fields of biomechanics (Panjabi et al., 1982a,b) and robotics (Karan and Vukobratovic, 1994). The problem is complicated due to the fact that human joints are not ideal, and therefore do not have a fixed centre of rotation. Lafortune et al. (1992) found that the joint centre of the knee compressed, moved front-to-back,

---

[14] http://www.naturalpoint.com/optitrack/. Last accessed 07/03/10.

and side-to-side by significant amounts during a normal walking cycle. In general, what is instead measured in biomechanics is the instantaneous Center of Rotation (COR) (Gerber and Matter, 1983), which is defined as the point of zero velocity during infinitesimally small motions of a rigid body.

To compute the instantaneous COR, markers are placed on each limb, and calculations made from the positions of the markers with the limb in different positions. The error in this calculation is reduced by using multiple markers for each joint, with a least squares fit used to filter redundant marker data (Challis, 1995). At the inception of this field, Spiegelman and Woo (1987) proposed a method for planar motions, which was then extended to general motion by Veldpaus et al. (1988). Due to the practical limitations on the number of markers that can be placed on a joint, further algorithms have been developed to optimise the placement of multiple markers (Crisco et al. (1994); Holzreiter (1991)). It is worth noting that there is an upper limit on the accuracy of joint estimations from motion capture data using this method, due to the aforementioned non-ideal human joints, and the fact that markers are attached to clothing or skin, not directly to the bone. Nigg and Herzog (1999) observed up to 3cm of skin movement over the tibia during ground running, which would directly translate into erronous joint estimations using the above method. This has lead to the development of other methods of computing human motion from collected motion capture data, principally the use of Inverse Kinematics (IK) to drive a virtual skeletal model (Molet et al., 1996; Boulic et al., 1998; Herda et al., 2000, 2001). Correctly modeling robust, accurate human motion from motion capture data, particularly from optical passive marker systems, is an area of current active research, with no single solution to the problems faced having yet been found.

1.3.3 Augmented Reality

AR in Stromohab presents novel issues, as conventional AR aims to fix virtual objects in static positions within a live, unpredictable video stream. Stromohab requires interactive (and therefore unpredictable) objects to be placed in a known video stream, an inversion of more standard requirements. AR in Stromohab therefore falls on a continuum between post-production techniques used in film-making (virtual objects placed, but with no interaction), and live AR methods currently demonstrated in augmented view applications (virtual objects static relative to real-world) such as (Starner et al., 1997; Vargas-Martin, 2002; Vargas-Martín et al., 2005; Ros et al., 2006; Peláez-Coca et al., 2011), Image registration is one of the most basic technical challenges in AR. Objects in the real and virtual environments must be properly aligned to ensure interaction is realistic. In VR applications, conflicts lead to disagreement in the visual-kinesthetic and visual proprioceptive systems, which may be a cause of motion sickness (Pausch et al., 1992), though because the kinesthetic and proprioceptive systems are less sensitive than the visual system (Burns et al., 2006), visual-kinesthetic and visual-proprioceptive conflicts are less noticeable than the visual-visual conflicts that occur in AR. There is a large body of literature in the field of AR (see Azuma and Others (1997) for a comprehensive explanation and review of the issues in the field), though the niche that Stromohab occupies appears to define a unique set of requirements, and therefore requires the development of of specialised technology to achieve the desired aims. The technical challenges presented by an Stromohab AR implementation do include some of the general challenges in the AR field, which are well described by Koller et al. (1997). For reference, these challenges are also briefly described here.

- Synchronisation of real-world and virtual cameras. A video-based AR system is created through the use of two cameras. A real camera generates video of the real environment, and a virtual camera, which generates the 3D graphics to be merged into the environment. To align virtual objects with the real environment, these cameras must have the same internal and external parameters.

- Precise description real object shape and location parameters. To allow accurate interactions between real and virtual objects, these parameters must be extracted from the real-world environment.

- Correct lighting. To generate convincing virtual objects, it is important to properly model the lighting of a real environment, and project it onto the virtual objects.

- Interaction with the environment. While this is stated as a primary challenge by Koller et al. (1997), the Stromohab project uses motion-capture to provide an intuitive interaction modality, and more generally, the recent advances in motion-capture technologies reduce the technical complexity of this area.

In summary, implementing AR for use in Stromohab presents a unique set of issues, most of which are however solved using a combination of established methods.

1.3.4 Usability and presence in a young population

Every year, the Electronics department at the University of York is involved in a science outreach event, in which teenagers aged 14-16 spend a day participating in a number of technology demonstrations, with the goal of inspiring them toward pursuing a career in engineering. Whilst no formal data was collected during the event, the participation of the Stromohab project provides some key insights into the required direction of the system development from a system perspective, and ultimately contributes towards focus of the formal experiments described in Chapters 3 and 4.

The demonstration consisted of a variation of the obstacle avoidance task described in Section 3.7. Students were secured into a fall-arrest harness, and then moved into position on a treadmill. An optical marker was attached to each of the students feet, and a corridor scene was presented on a portrait monitor in front of the treadmill. The position and motion of the optical markers was represented by three-dimensional representations of 'feet' within the virtual environment. On the floor of the virtual corridor scene were a number of three-dimensional boxes. Each box was decorated with one of two recognisable public figures, and the students were instructed to either 'kick', or 'avoid' the boxes, based on the decoration. As the treadmill speed was increased to up to a slow walking pace of 2 mph, the virtual corridor optic flow increased at the same rate.

While individual performances varied within the task, overall a number of consistent observations were made.

1. Immediate sense of presence and engagement. No student had great difficulty in grasping the control paradigm. The use of motion-capture with real time visual feedback is intuitive, and after a few initial movements, with minimal instruction, students were confident in their control of their representation onscreen.

2. Difficulty in depth judgement in the task. While the students were able to control their avatar easily, there appeared to be an issue with their ability to judge their position 'into' the scene. This was evidenced in their collisions with the virtual boxes. While there were collisions due an incorrect judgement of the width or the height of the boxes, in the majority of cases the collisions occurred when an attempt to step 'over' a virtual box was made. In general, the students would lower their foot too early, in essence making their step too short to completely avoid a box.

These observations drive the quantification of the effect described in Section 3.2, and the design of the complete virtual task presented in Section 3.7.

1.4 Summary

Thus far, stroke, and its acute and long-term effects have been described (Section 1.2). The current primary aims and approaches of stroke rehabilitation have been discussed and evaluated (Section 1.2.1), and the motivations for utilising virtual-reality and motion-capture technologies have been presented (Section 1.3). Virtual environments and their application to rehabilitation (Section 1.3.1) have both been explored, followed by a detailed review of the technologies available for human motion tracking (Section 1.3.2). Finally, brief explorations of human-motion reconstruction (Section 1.3.2), and the technical challenges presented by the use an augmented-reality virtual environment of the type proposed for use in the Stromohab project (Section 1.3.3) are described.

Chapter 2 presents the design and implementation of the Stromohab research platform, as driven by the above review and further exploration of the literature. I present a series of experimental studies conducted using the Stromohab platform (Chapter 3), in which a subset of the capabilities of the platform are demonstrated through the use of the system to: quantify the problem of distance underestimation in virtual environments (Sections 3.2), investigate a resolution to this problem (Sections 3.3 – 3.6.3), and confirm that the application of a proposed model (Section 4.5) resolves the issue (Section 3.7).

Discussion of the work presented by this thesis can be found in Chapter 4, which includes: a discussion of the experimental results of each study, critical review of the work, and an exploration of potential future directions for this research.

Chapter 2

Technical System

Chapter 1 shows that there are very few post-stroke VR-based rehabilitation systems currently targeting lower body and gait rehabilitation, and previous work on the Stromohab prototype system formalised functional requirements and went some way to investigating and implementing solutions to the technical challenges involved (Dring, 2007; Holmes, 2007; Grimes, 2008). However, while the results of this work had produced an initial proof of concept system, it also demonstrated that there were technical limitations to the initial design. Furthermore, investigations into the type of clinical environment the system was likely to be used in, and the availability of new hardware, led to a new, improved system design and implementation.

In this chapter, I use established engineering design methodologies to map out the problem space, systematically analysing and specifying desirable properties and features for the system, before implementing and subsequently evaluating it against these objectives. As the current state of the art does not clearly describe all of the intrinsically desirable characteristics of a system such as the Stromohab platform, the final design is approached iteratively, building on a number of case studies. Thus, a problem that is initially *open-ended* and *ill-defined* (Dandy, 2007) becomes well-framed, with objectives that are both deliverable and testable.

The process of this design and implementation is described below.

2.1 Engineering Methodologies

The two current prevailing software engineering development methodologies are the Agile and Waterfall methods of software design. While the Agile method is becoming increasing adopted, (Sumrell, 2007; Sureshchandra and Shrinivasavadhani, 2008; Desouza, 2012), there are numerous discussions and proponents of both methodologies (Huo et al., 2004; Sliger, 2006; West et al., 2010), and as such, a small discussion is warranted here.

Waterfall

The Waterfall software development model is classically linear and sequential, in which each waterfall stage in assigned to a separate team to ensure greater project and delivery deadline control. The linear approach includes the following (high-level) stages:

1. Analysis: The project team first analyses the problem, and then determines and prioritises the business requirements and needs.
2. Design: The requirements are then translated into IT solutions, and decisions made regarding the underlying languages and technologies are made.
3. Implementation: Programming code implementing the completed design is then produced.
4. Evaluation and Maintenance: The completed project is evaluated against the original requirements, deployed into production and maintained as required. Should any issues arise in this stage, generally either implementation bugs arising from incorrectly written code, or incorrect functionality arising from an erroneous problem analysis, the entire process must be repeated.

Agile

The Agile software development model is a low-overhead method that emphasizes values and principles over processes. Working within cycles typically consisting of a week, fortnight or month in length, project priorities are re-evaluated at the end of each cycle. The Agile methodologies guiding principles can be summarised as valuing:

1. Individuals and interactions over processes and tools.

2. Working software over comprehensive documentation.
3. Customer collaboration over contract negotiation.
4. Responding to change over following an initial plan verbatim.

Comparison

To summarise the main difference between the methodologies, Waterfall can be viewed as requiring predictability, while the Agile method values adaptability. Agile methods reduce project overheads such as rationale, documentation and meetings, and as such, benefit small software teams delivering on requirements that are likely to change. Waterfall relies on the strict enforcement of methods and processes, while the Agile method provides a much smaller feedback loop. Agile breaks a large overall project down into small, functional pieces, with each piece designed and implemented separately. This minimises the overhead in adjusting to new requirements, or correcting issues from the design or implementation process.

As the Stromohab system is intended for use within a research environment, by definition the features and functionality required cannot all be initially defined. Therefore, the primary need for flexibility and adaptability have lead to the adoption of the Agile method in the design of the Stromohab system.

2.2 System Overview

At the conclusion of the project, the Stromohab software consists of approximately 180,000 lines of primarily C# code. It consists of three main executable modules, each capable of communicating over any TCP/IP network. A reliable, robust public API has been developed, and tested, and has already been used by multiple third parties who have worked with the Stromohab system. The system is fully documented, with HTML documentation accessed via a web server. All code is under revision control, providing both simple access and efficient merging of features written by third parties. In addition to the main system, a MySQL database provides a reliable and relational backing data store, accessed using a custom data access layer. This DAL provides database independence, allowing for future development and expansion as required. The system provides a controllable, engaging virtual environment for a patient or test subject, and an intuitive user interface for a researcher or therapist.

While this thesis will not cover the design of the system in detail, an overview of the system components can be found in Figure 2.1.

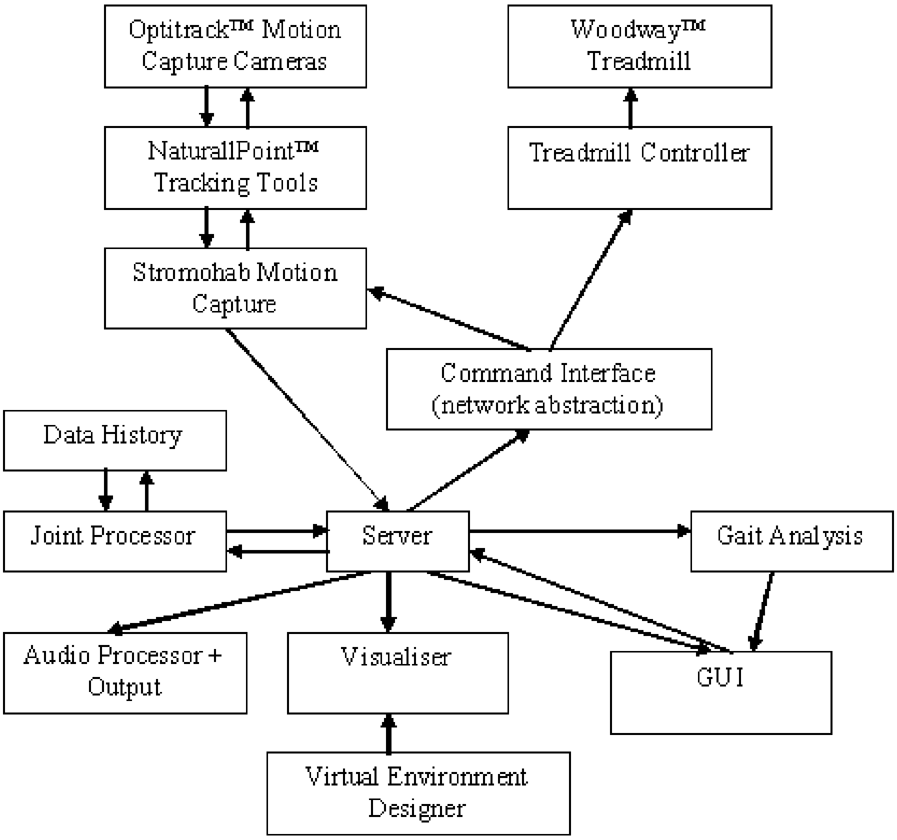

Figure 2.1: *Current Stromohab system structure*

The main system modules can be broadly collated into the following :

1. Hardware Interface and Network Abstraction (*TCP Server*[1]).
2. Control Interface (*GUI*).
3. Interactive Graphical Display (*Visualiser*).

Each of these modules runs as an encapsulated executable, controlling a number of the subsystems shown in Figure 2.1.

---

[1] Transmission Control Protocol: The transport layer of TCP/IP networking. In this context, TCP Server is used to describe the Stromohab Server application - a custom network server built around raw TCP sockets.

2.2.1 TCP Server

The motion capture hardware currently consists of ten Optitrack™ V100:R2 cameras, hardware synchronised and connected to a Microsoft Windows® PC via two dedicated high speed USB hubs. Once calibrated, the Optitrack™ motion capture system provides two routes to accessing marker position data. The system comes supplied with Arena™ studio software: this provides a data collection workflow suited to the collection of motion capture data for post-processing and subsequent character animation. The Stromohab system however, requires access to the stream of marker data in real time, and as such it uses an alternate data collection capability. This comes in the form of a software library supplied with the motion capture system, which provides low level access from bespoke software such as Stromohab, via dynamic linking.

The TCP Server wraps this low-level library and provides transparent network access to the library functionality. It handles code locking and eliminates race conditions, allowing multiple clients to connect via TCP and control the motion capture hardware. Additionally, it provides multiple, synchronised real-time data streams to an essentially unlimited number of connected clients, via a custom low-level transfer protocol. This then allows, for example, the presentation of a real-time virtual environment to the patient or subject, in conjunction with the presentation of a real-time control and diagnostics interface to researchers and clinicians. It also allows the connection of numerous parallel research clients, a feature that proved invaluable during the development of the system, when multiple researcher often required access to the system at the same time.

As all of the data is available via a transparent network interface, remote control and monitoring of the system is possible, providing clinicians the option of controlling and monitoring physiotherapy sessions remotely, as has been presented as desirable in prior literature (Deutsch et al. (2001); Boian et al. (2003) present non real-time access to patient data, Lai et al. (2004) present a telerehabilitation model for community-based stroke rehabilitation, and Tognetti et al. (2005) produce a wearable system aiming to allow remote treatment of post-stroke patients).

Further to this, through the use of UDP packet broadcasting, the server identifies itself to any connected network once every ten seconds. This allows clients to connect to the server, even when they have no prior knowledge of the server's location or technical set up. This in turn provides portability to the Stromohab system, allowing it to be moved, disconnected from a network, and reconnected elsewhere, with no change in configuration required for non-technical users of the systems - i.e. clinicians, physiotherapists, and researchers utilising only the endpoint of the positional data stream.

Via the Stromohab motion capture subsystem, the TCP server also performs a a number of additional clean-up operations on the raw camera data stream. It enhances and extends the functionality of the Optitrack™ software library, performing operations such as marker list management, which prunes markers that are no longer visible to the cameras from the data stream. It also provides options to separate, group and trajectorize raw marker data, such that a marker on the left foot, for example, is consistently tracked as distinct from a marker on the right foot (see Section 1.3.2 for a description of the marker identification limitations of passive optical motion capture systems).

Access to the functionality of the TCP Server is provided by a well-documented API, bundled into a dynamically linked software library that is provided to developers of client software. This allows new researchers and developers to rapidly access and use the Stromohab system, in part by hiding the complexities of parallel raw data stream access and management.

Treadmill Controller

The Stromohab treadmill controller is an interface specifically developed for use within the LIVE environment at the University of York. Our laboratory treadmill is a Woodway™ model that includes a serial port communication interface. Through the use of a custom developed software driver, the treadmill controller provides various options to control the operation of the treadmill from the PC. These include starting and stopping the treadmill, increasing and decreasing forward and reverse speeds, and adjusting the incline of the treadmill. This capability allows for complex virtual environments to be created, in which hills, for example, are simulated through the adjustment of the treadmill incline. It also allows for a single researcher to precisely control and direct experimental sessions via the GUI, ensuring that variables such as movement through a virtual environment is precisely aligned with the movement of the treadmill belt.

2.2.2 Graphical User Interface

The Stromohab graphical user interface has been through a number of design iterations, resulting in a number of control interfaces, each suited for differing environments. The initial user interface is a Windows® forms based C# application, communicating with the motion capture hardware via the TCP

Server described above. This interface is designed primarily for use within a research environment, and as such, focuses on allowing the user to create and run experiments, storing data in an easy to analyse flat-file format. Forms allowing the creation, and loading and saving of virtual tasks, the recording of subject details, and the scripted control of an experimental session are provided, producing an encapsulated experimental set up that aims to maximise subject throughput. The virtual environment designer allows a researcher to rapidly create new virtual environments and tasks, and communicates in real-time with the Visualiser, via a custom TCP protocol. This allows the rapid design and testing of new virtual tasks, as any changes to the task are immediately reflected on screen, and tasks may be trialled at any point in the design process.

Gait Analysis

During an experimental session, in additional to any experiment specific measures, the GUI displays a number of real-time spatio-temporal gait metrics, such as current stride length, average stride length, stride duration, foot velocity and gait asymmetry index. Real-time graphs are also displayed, indicating the position of the foot in the transverse, coronal and sagittal planes. These allow the researcher to compare the paths of the left and right feet, and immediately see any gait abnormalities.

In addition to metrics reliant on only the foot position, the system also allows for more advanced gait analysis, and computes limb positions and angles, such as the current, maximum and minimum dorsiflexion angle. Metrics such as stride length and stride frequency are calculated in real-time via continuous processing of a marker position. Sensor noise is first smoothed from the position data using a moving average temporal filter, after which turning points are calculated by assessment of the coronal and sagittal paths of the marker.

To allow the animation of a full-body avatar, subject joint positions are calculated in real-time using a COR technique, as described in Section 1.3.2.

GUI Version II

The second iteration of the Stromohab GUI focuses on the development of an interface suitable for use within a clinical environment. This interface was created based on the Windows® Presentation Foundation (WPF) framework, which provided a more modern look and feel to the graphical interface, as well as allowing an intuitive workflow to be implemented. Figure 2.2 shows the overall implemented GUI workflow. It was designed to be as straightforward as possible, and to eliminate as many unecessary steps as possible.

GUI Version III

The final iteration of the GUI was implemented based on feedback from clinicians and physical therapists based at Addenbrooke's Hospital, Cambridge. Key changes from earlier versions include a further streamlining of the clinical workflow, and an adaptation of the GUI such that it is suitable for use via a touch-screen display, as opposed to the previous, standard keyboard and mouse control paradigms.

2.2.3 Visualiser

The Stromohab Visualiser has also been through a number of design iterations. The first version was developed for use in a research setting, and consists of an SDL.NET based graphical environment, running in a maximised, windowed mode. This virtual environment is completely hand-coded, and as such provides extremely low-level design and control to the researcher. Graphical objects are implemented using OpenGL primitives, with additional detail being provided through the use of textures.

One interesting problem often present when designing virtual environments for use with motion capture is that of footskate (Kovar et al., 2002; Zordan and Van Der Horst, 2003; Zordan et al., 2005). This is the problem of avatar foot representations appearing to 'slide' across a virtual environment floor, due to differences in the visual flow of a background and the captured motion of a subject. There have been a number of attempts to quantify and resolve this issue, but at the time of writing, no complete solution has been found. The Stromohab experimental setup however, presents a special case in terms of footskate, as the treadmill belt provides a surface moving at the speed required within the virtual environment.

Taking advantage of this, highly reflective optical markers, in the form of stickers, are attached to the treadmill belt at semi-regular intervals. These intervals are spaced such that at least one reflective marker is visible to the motion capture cameras at any one time. The Stromohab motion capture subsystem, as seen in Figure 2.1, then captures these markers, calculating the current treadmill belt speed and notifying the Visualiser (via the TCP Server) of this speed. This approach provides a number of advantages over a manual set up. In a purely

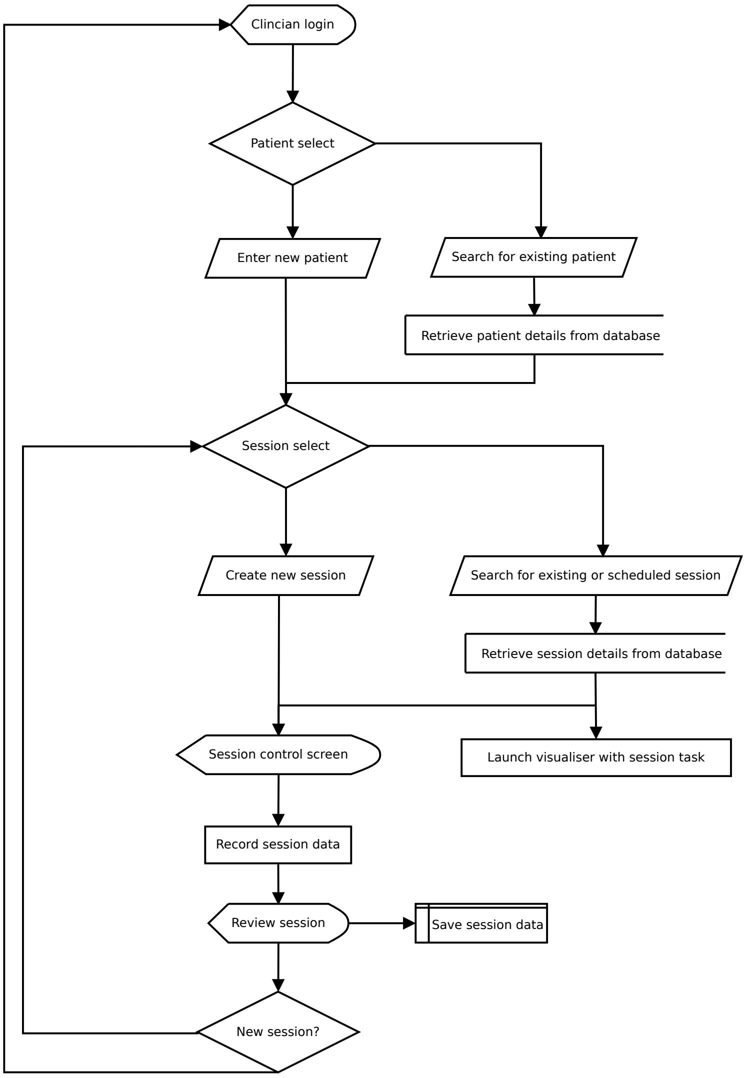

Figure 2.2: *Clinician Interface workflow*

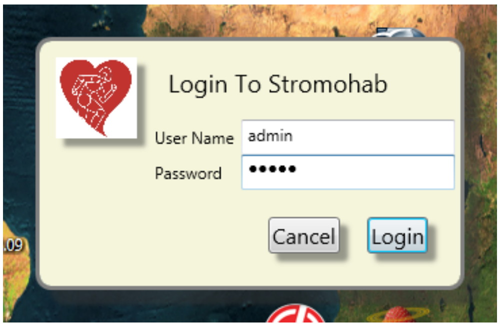

User Name admin

Password •••••

(a) *Login screen*

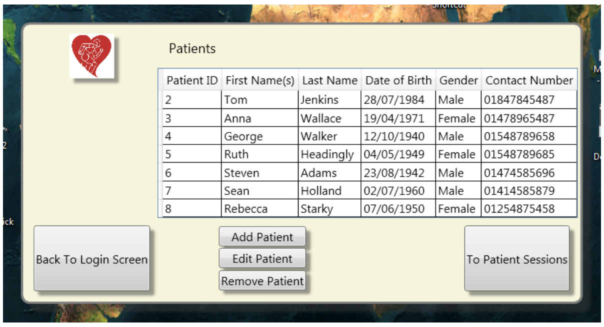

Patients

Patient ID	First Name(s)	Last Name	Date of Birth	Gender	Contact Number
2	Tom	Jenkins	28/07/1984	Male	01847845487
3	Anna	Wallace	19/04/1971	Female	01478965487
4	George	Walker	12/10/1940	Male	01548789658
5	Ruth	Headingly	04/05/1949	Female	01548789685
6	Steven	Adams	23/08/1942	Male	01474585696
7	Sean	Holland	02/07/1960	Male	01414585879
8	Rebecca	Starky	07/06/1950	Female	01254875458

(b) *Patient list*

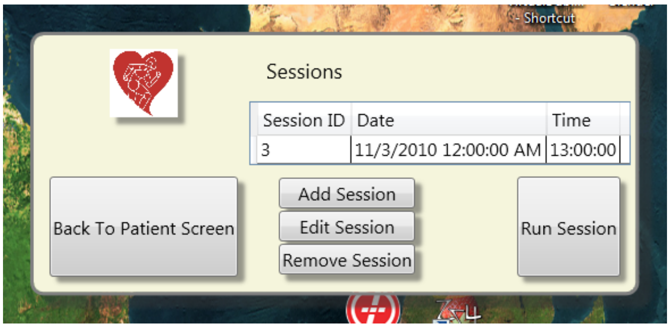

Sessions

Session ID	Date	Time
3	11/3/2010 12:00:00 AM	13:00:00

(c) *Rehabilitation session list*

Figure 2.3: *The streamlined Stromohab GUI workflow. Feedback from clincial staff was highly positive, as a rehabilitation session can be launched in only five clicks from system starup.*

```
(run for 30 seconds)

(calibration)
for each detected marker
  log minimum and maximum x values
  log minimum and maximum y values

(filtering)
for each detected marker
  is marker within x range ±5mm?
    is marker within y range ±15mm?
      calculate marker (and therefore belt) speed based on current and previous positions
      remove marker from processing pipeline
```

Listing 2.1: *Treadmill calibration routine*

manual set up, though it is possible to set the speed of the virtual environment optic flow, and then set the treadmill to a similar speed, significant visuomotor conflict is experienced until the two separate systems align.

The capturing and subsequent filtering of the treadmill markers is fully automated, following an initial calibration step. After attaching the reflective marker stickers to the treadmill belt, a the treadmill is started, and the belt speed increased to a speed close 1mph. The calibration routine is then started via a button press on the GUI. This calibration and filtering routine is summarised in Listing 2.1.

This routine has the effect of isolating a narrow ‘corridor’ of expected treadmill speed markers, and allows highly efficient detection and processing of the treadmill belt speed, while also eliminating additional marker identification and processing further along the processing pipeline.

An alternate method of detecting the speed of a subject moving through a virtual environment is described in Listing 2.2. This approach more closely approximates the experience of walking in the real world, in which progress is not made at a consistent smooth rate, but in a more sinusoidal pattern, acceleration and deceleration occurring within each complete stride. After testing however, this approach was found to ‘feel’ very discontinuous, and was deemed to negatively affect engagement and realism. It is therefore not used in the current Stromohab system.

Illumination Model

Related to the issue of footskate, within virtual environments it can be difficult to judge the height at which an avatar, or virtual object is positioned, relative

```
(for each frame)
identify each subject foot marker
detect foot down event and subsequent stance motion
progress visualiser forward a distance equalivalent and opposite to
  that moved by the foot
```

Listing 2.2: *Alternate speed detection routine*

to the floor. In complex graphical environments, this issue is minimised through the use of complex lighting models, applying effects such as soft shadows, specular highlighting, and diffuse shading. One of the aims of the initial Stromohab visualiser was to minimise the complexity of the virtual environment, and as such, the lighting model only included a diffuse, global lighting model. While this allows the scene to be lit, and ensures objects are visible and coloured, it also does not include calculations for shadowing, etc. Therefore, a simple shadowing model is applied to the avatar, with which a proportional overhead projection of the avatar shape is projected onto the ground-plane, beneath the avatar model. The size and shape of this projection is calculated using simple scaling factor. Each for each visual frame, the X (width) and Z (depth) dimensions are scaled by a percentage m.

Figure 2.4 illustrates this illumination model. An overhead light source is modelled, which proceeds through the virtual environment with the same velocity as the avatar. If this light source is positioned at position d above the ground-plane, then $m = 0$ when the avatar is at position $Y = d/2$, and $m = 100$ when the avatar is at position $Y = 0$ and $Y = d$.

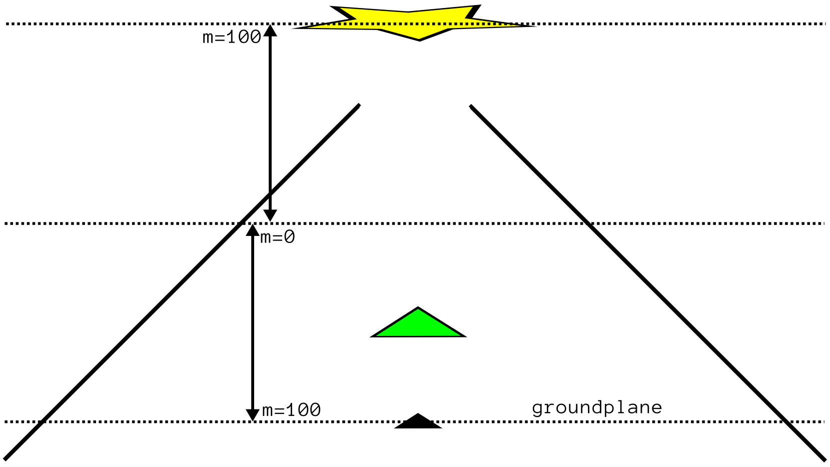

Figure 2.4: *Stromohab Visualiser illumination model version I*

Collision Detection

Thus far, we have created a virtual environment which allows the real-time motion-captured movement of a subject to be represented on-screen, within an a graphical virtual environment. While this is useful, the power of virtual environments comes from the interactive experience that it is possible to create with the paradigm. Within virtual environments, interactions between objects are implemented utilising a *collision detection* technique. While on-screen collision detection techniques have been developed and implemented for a number of decades, (Boyse, 1979; Hirzinger et al., 1989; Vaněkčkek Jr, 1994), efficient methods are still under active investigation. Teschner et al. (2005) for example, provide an in-depth review of collision detection for deformable objects, and Xu et al. (2007) present an efficient collision detection method applicable to robotic bridge maintenance

To explore the collision detection algorithm used in the Stromohab Visualiser, we will first examine a simple two-dimensional example, before expanding the concept to include three dimensions.

In the simplest case, we have two objects, A and B. Let A have dimensions $x2-x1$ and $y2-y1$, as seen in Figure 2.5. Should we require interaction between objects A and B, based on their status (*colliding / non-colliding*), we are interested in knowing if the objects are overlapping, as seen in Figure 2.6. Though there are many approaches that can be taken to solve this problem, including assessing the state of a pixel before drawing over it, an efficient scaling approach can be implemented using a *bounding box* methodology. In this example, we can assess each of the vertices of B, and consider if they fall within the boundaries of A. Our algorithm can thus be summarised as shown in Listing 2.3.

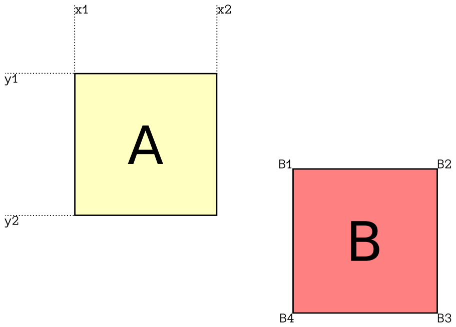

Figure 2.5: *Collision detection target objects*

```
for each vertices V in B
if (V(x) > x1(A) and V(x) < x2(A) ) {
```

```
3    if (V(y) > y1(A) and V(y) < y2(A) ) {
4      A and B have collided
5  }
6  else {
7    no collision
8  }
```

Listing 2.3: *Simple two-dimensional collision detection*

Figure 2.7 illustrates the state of the algorithm after Line 2, before we have confirmed that the objects are overlapping in both the X and Y dimensions as in Figure 2.6.

Once we have implemented this algorithm, we can therefore test for collisions between any two objects in our environment by iterating through all objects within the environment and testing using the above algorithm. As the number of potential object interactions increases, this approach rapidly becomes inefficient. In the Stromohab case however, collisions in most virtual tasks are constrained to interactions between the avatar and the environmental objects, leading to an efficiency of $O(N + 2)$, where N is the number of objects in the virtual scene, and the potential number of avatar interaction objects is two.

Visualiser Version II

An alternate version of the virtual reality visualiser was implemented using a managed version of the *OGRE* 3D games environment (*MOGRE*[2]). This environment was implemented to aid in the rapid creation of new virtual environments, and to allow more complex environmental interactions and physics. Using an established game engine to render the graphical environments also allows non-technical users of the Stromohab system to create new virtual environ-

---

[2] `http://www.ogre3d.org/tikiwiki/tiki-index.php?page=MOGRE`

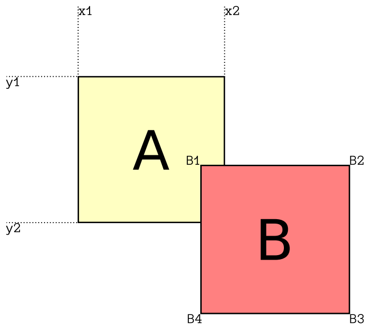

Figure 2.6: *Collision detection target objects overlapping*

ments. While any graphical editor that supports export to the *dotScene* format may be used to create virtual environments, the primary editor recommended for use with Stromohab is the open-source *Ogitor SceneBuilder*[3] software. This software provides an intuitive user interface, and aids in the rapid creation of virtual tasks and environments of the type most often required for use with the Stromohab platform. Figure 2.8 shows an example of the *Ogitor* editor.

Additionally, in order to easily support more advanced collision and overall physics in this version, the Nvidia™PhysX™engine is utilised, via the managed C# bindings in the eyecm-PhysX Candy Wrapper[4]. The PhysX™library provides highly optimised implementations of both collision detection and physics simulations, and also allows the processor intensive operations required for these complex simulations to be offloaded to dedicated hardware on a Nvidea™graphics card. While initially more complex to implement, the integration of this library into the Stromohab platform now allows advanced physics and other features to easily be included into virtual environments and tasks. This in turn provides benefits in terms of realism and engagement, further enhancing the potential benefits of the platform in both research and clinical environments.

Augmented Reality Visualiser

A video capture virtual environment places a live video feed of the subject inside a virtual environment, and the use of video based virtual environments has been suggested to add to the realism and the users sense of presence in a virtual environment (Nash et al., 2000). Rand et al. (2005) find this to be true, when a video capture virtual environment is used. In an attempt to take this idea further, the concept of an augmented reality environment was conceived and

---

[3] `http://www.ogitor.org/`

[4] `http://eyecm-physx.sourceforge.net/`. Last accessed 15/12/2012.

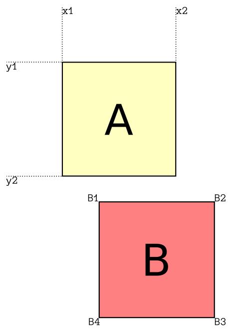

Figure 2.7: *Collision detection in X*

implemented. This is distinct from current literature describing video-based virtual environments - a video-based virtual environment places a live view of a subject inside a computer generated virtual environment, while our concept of an augmented reality virtual environment takes an avatar representation of a subject and places them within a pre-recorded video environment. For example, the pre-recorded video environment might consisted of a walk along a countryside path, or to further enhance a subject's sense of presence, a familiar scene may be captured and turned into an augmented environment. This scene may then be augmented with both a virtual avatar, and virtual objects, in order to create a useful virtual task. A sample augmented reality obstacle avoidance task can be seen in Figure 2.9.

In order to implement an augmented-reality environment of this type, a number of technical challenges must be overcome. The primary, and perhaps most most difficult of these is found in the initial video capture of the environment. The difficulty lies in recording a stable motion through a scene of great enough length to be useful as an environment. As an example, a capturing a video environment that allows five minutes of forward motion when utilised in a virtual task requires a stable video capture of at least five minutes. During this video capture, the recording camera must move through the scene at a constant speed, towards a constant focal point. Oscillatory camera motion must essentially be non-existent, as any camera 'shake' diminishes the usefulness and sense of presence when the video is later used as a virtual environment.

There are few examples of this type of shot in classical film productions, and none of the length required for use as a virtual environment. To achieve a long, smooth shot, videographers often use a 'dolly' - a section of smooth track, along which a camera and operator are wheeled. This type of setup however, will generally produce panning shots, in which the camera sweeps across a scene horizontally, rather than the motion 'into' the scene required here, and the amount of track required to produce a continuous shot of the length required in this application is highly prohibitive. Therefore, a number of approaches have been trialled during the Stromohab project, including the use of a hand-held and fitted Steadycam rig, in which the video camera is counterweighted to minimise unwanted motion. Other approaches have included mounting the camera on wheeled vehicles, shooting video from palette trolleys as well as from bicycles and cars. These approaches each solve, and create their own issues. Principally, methods involving a camera supported by the body, directly, or via a Steadycam setup, lead to low frequency oscillations in the final captured video, whereas methods in which the camera is affixed to a wheeled vehicle produce high frequency oscillations. In essence, without the large investment

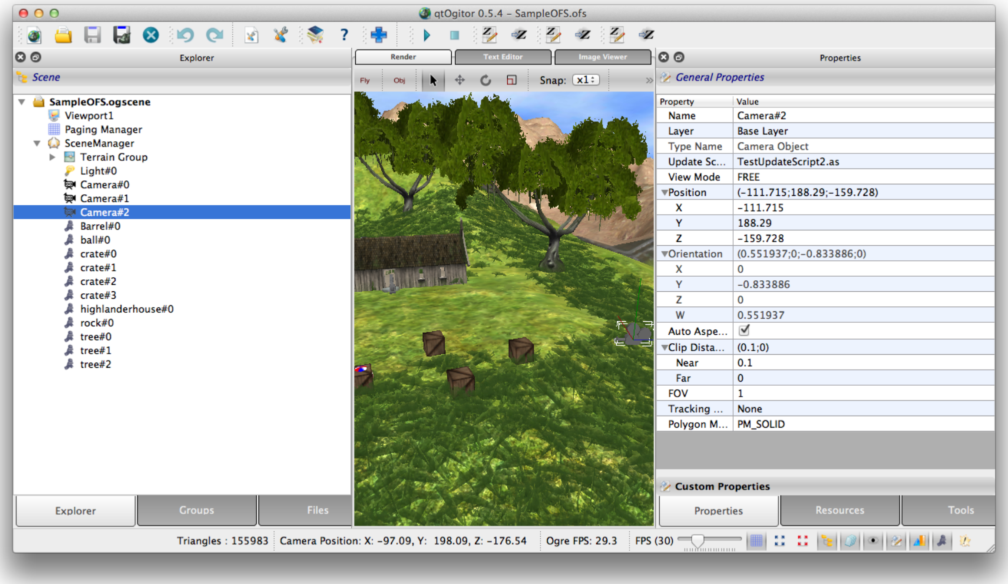

- SampleOFS.ogscene
  - Viewport1
  - Paging Manager
  - SceneManager
    - Terrain Group
    - Light#0
    - Camera#0
    - Camera#1
    - Camera#2
    - Barrel#0
    - ball#0
    - crate#0
    - crate#1
    - crate#2
    - crate#3
    - highlanderhouse#0
    - rock#0
    - tree#0
    - tree#1
    - tree#2

Property	Value
Name	Camera#2
Layer	Base Layer
Type Name	Camera Object
Update Sc...	TestUpdateScript2.as
View Mode	FREE
Position	(-111.715;188.29;-159.728)
X	-111.715
Y	188.29
Z	-159.728
Orientation	(0.551937;0;-0.833886;0)
X	0
Y	-0.833886
Z	0
W	0.551937
Auto Aspe...	☑
Clip Dista...	(0.1;0)
Near	0.1
Far	0
FOV	1
Tracking ...	None
Polygon M...	PM_SOLID

Triangles : 155983 Camera Position: X: -97.09, Y: 198.09, Z: -176.54 Ogre FPS: 29.3 FPS (30)

Figure 2.8: *Ogitor SceneBuilder user interface. From* `http://tracker.ogitor.org/uploads/857/ScreenShot2012-04-19at9.42.03PM.png`. *Last accessed 24/10/2012.*

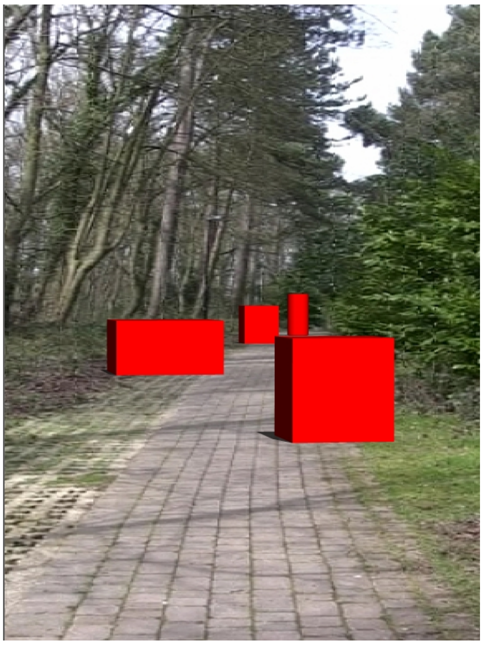

Figure 2.9: *Augmented reality obstacle avoidance sample task*

of film quality production equipment, it is difficult to capture the type of video required for use in an augmented reality environment.

The issues with capturing suitable, stable video therefore logically lead to a post-production solution. There are a number of approaches to post-production video stabilisation, with each again coming with their own set of advantages and disadvantages.

The first approach we will examine here is that of filtering. This approach attempts to minimise image change between frames, using frame interpolation to 'fill in' image segments removed by the filtering. One of the popular open source video filter applications is the DeShaker filter[5] for the VirtualDub[6] video capture/processing utility. This filter is a two-pass filter, and operates as follows. During the first pass, the filter attempts to find the panning, rotation and zoom that will make the current frame look as similar to the previous frame as possible. It achieves this by sampling pixel blocks in each frame, then finding the optimal matching shift. Large motions are removed by scaling down the images and calculating an initial shift from the whole (scaled) images. This initial shift is then optimised by doubling the size of the image and processing four blocks, and repeating this process until the block size is equivalent to a user-defined value. The motion vectors for each frame are thus derived, and from these, the optimal values for panning, rotation and zoom are calculated and written to a log file. The second pass then uses these values to calculate an optimal, simulated camera motion based on user-defined smoothness settings. Each video frame is then transformed based on this motion.

When the VirtualDub Deshaker approach was tested, it was found to have a number of limitations. While the filter has a large number of user-definable parameters, we found that it was not possible to stabilise the video to the extent required, without introducing an excessive amount of distortion. This distortion is an artefact primarily introduced by the algorithm interpolation mechanism, and is proportional to the amount of smoothing required. It is possible to minimise the distortion by focusing the filter on only eliminating either high or low frequency camera motion; however our testing concluded that this approach did not produce smoothed video of a quality high enough for use in a video based augmented reality environment.

A second approach to video stabilisation is that of video camera tracking, or match-moving. This is the process of extracting the real-world motion of a camera from video, and is generally used by the film industry to add computer

---

[5] `http://www.guthspot.se/video/deshaker.htm`. Last accessed 07/01/13.

[6] `http://www.virtualdub.org/`. Last accessed 07/01/13.

generated special effects into pre-recorded video.

Matching-moving can be broken down into two major steps. The first is feature identification and tracking. A feature is a specific point in the video that can be isolated and followed through multiple frames. There are a number of frequently used feature tracking algorithms, generally based on colour, or edge detection image processing techniques. As a feature is tracked across multiple frames, it is transformed into a series of two-dimensional screen coordinates known as a track. The second step then attempts to derive the three-dimensional motion of the camera by solving the inverse projection of a complete set of tracks.

This was the approach taken to stabilise the captured video used as virtual environment in the Stromohab AR Visualiser. The open-source Voodoo[7] camera tracking was initially explored for this process, however software stability proved to be a significant issue and so the popular SynthEyes™ camera tracking and stabilisation commercial software application[8] was used.

2.2.4 Physical System Development

The physical hardware that makes up the Stromohab system has also experienced a number of design iterations. After the initial system concept had been prototyped and proven using the motion-capture hardware described in Section 2.2.1, a lightweight frame was designed and constructed to integrate the various hardware components into a complete, coherent system. The primary design features of the frame include: a wheeled base, maximising portability of the system and allowing its use both on and off the system; a powered extensible vertical support, allowing adjustment of the display screen height to suit different purposes; a custom-engineered display mount, allowing tilting of the display screen up to 60° from vertical (Casey et al., 2010; Pelah, 2007); a lightweight tubular framework, providing multiple motion-capture camera mounting options.

Figure 2.10 shows the display mount mid-construction, with the tilting mechanism partially installed. This mechanism is precisely positioned to rotate the heavy display screen exactly around its centre of gravity, allowing adjustment of the screen tilt with minimal force.

Figure 2.11 shows the fully assembled system, in situ while exhibiting at Venturefest Yorkshire 2010. This exhibition provided us with an ideal opportunity

---

[7] `http://www.digilab.uni-hannover.de/docs/manual.html`. Last accessed 18/12/12.

[8] `http://www.ssontech.com/index.html`. Last accessed 18/12/12.

to test the reliability and robustness of the system and its performance within a limited space. This successful field testing demonstrated the potential for the Stromohab system to move beyond a laboratory setting, toward a clinical deployment.

The second iteration of the Stromohab system directly aimed at producing a system ready for clinical deployment. While this system required all of the functionality of the previous iteration, additional constraints included the requirement to encapsulate all physical hardware, while remaining easy to maneuver by a single person. A number of designs were considered, including those illustrated in Figure 2.13. A variation on Figure 2.13b was assessed to provide the most flexibility, and concept was manufactured. The complete Stromohab clinical system can be seen in Figure 2.14.

To conclude this whistlestop tour through the technical design and implementation of the Stromohab research and rehabilitation platform, we note the incremental design improvements, with an emphasis on field testing and feedback from clinical environments. One motivation for this is further discussed in Section 4.7.1, in which it is noted that VR-based rehabilitation systems and technologies are only just beginning to filter through to the clinical trial stage. A motivation of this thesis is the desire to encourage more interdisciplinary collaboration in this field, in order to best utilise potentially beneficial technologies as effectively as possible.

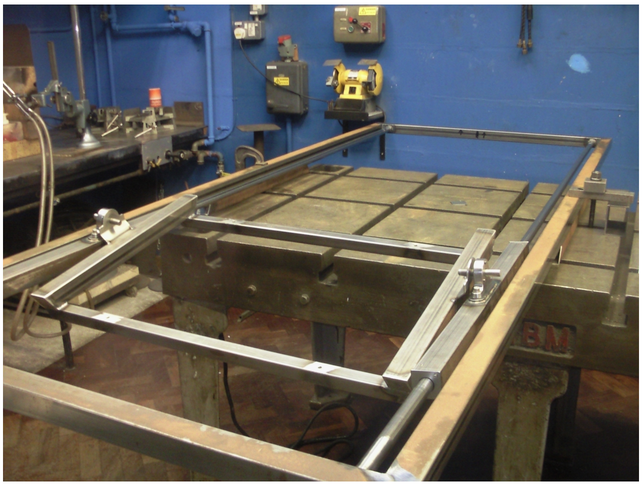

Figure 2.10: *The custom-engineered display screen mount, under construction. The frame is capable of supporting a 50 inch plasma television mounted in a portrait configuration, weighing up to 150kg.*

Figure 2.11: *The fully assembled initial system, shown here in front of a portable treadmill at while exhibiting at Venturefest Yorkshire 2010. The virtual corridor can be seen on the screen, positioned at the beginning of an obstacle avoidance task (Section 3.7).*

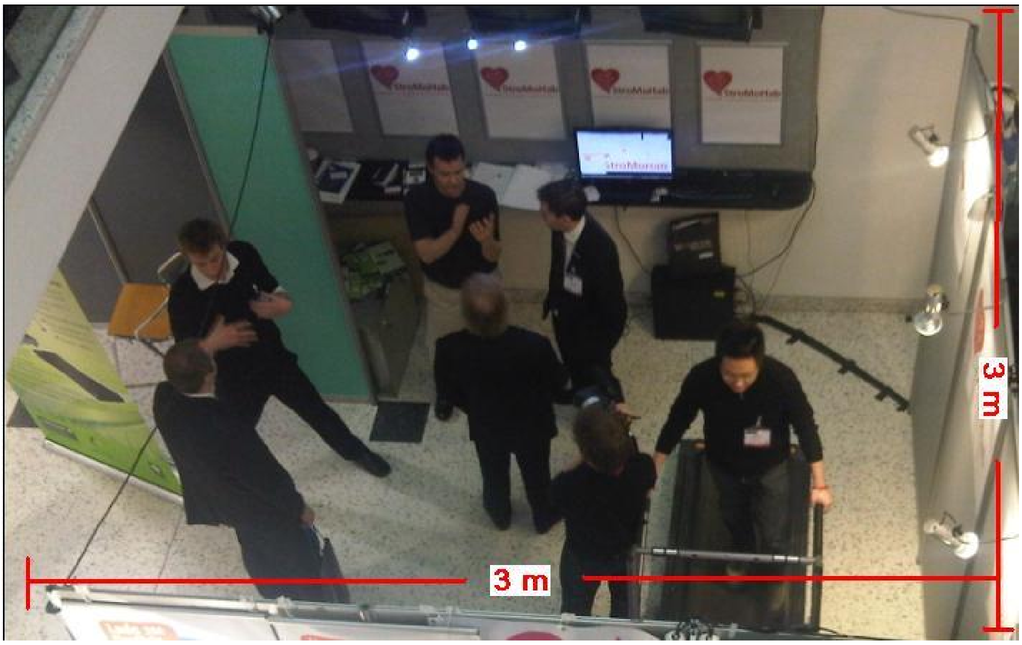

Figure 2.12: *The Stromohab exhibition booth at Venturefest Yorkshire 2010.*

(a) *Fixed cameras* **(b)** *Wide cameras*

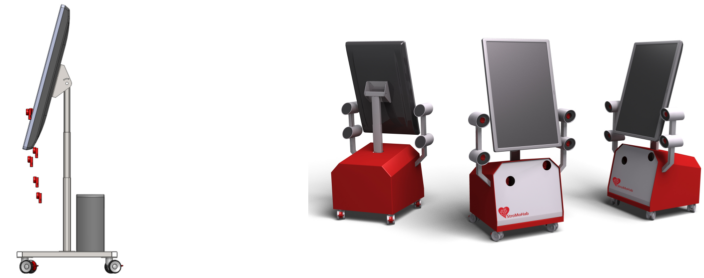

(c) *External cameras* **(d)** *Pod cameras*

Figure 2.13: *Stromohab clinical system design concepts.*

Figure 2.14: *Stromohab clinical system, the second physical design iteration.*

Chapter 3

Psychophysical Experiments

The research presented in this thesis spans multiple disciplines and areas of study, and broadly speaking, these can be broken down into two distinct areas - engineering design and psychophysical research. This chapter describes the psychophysical experiments. The general rationale for the experiments will be presented, followed by a general explanation of the common methodology used throughout the psychophysical research. Finally, the specific hypotheses, methods, results and discussions of each experiment will be presented in turn.

There were two general aims that motivated the following set of psychophysical experiments. The first aim was to further test and demonstrate the potential applicability of the Stromohab system. These experiments were conducted on healthy controls in a research setting, with the understanding that there are numerous additional factors to be considered when implementing such a system in a clinical rehabilitation environment. While the Stromohab system was designed for eventual use with stroke patients in a clinical rehabilitation setting, the system must first prove to be reliable in a controlled environment. The following set of experiments will show that Stromohab is able to measure and record subject movement with the degree of accuracy required for use as a dependent variable in the detection of differences in motor performance.

A second general aim was to identify the key factors that affect performance on motion-capture motor tasks that necessitate the representation of the user's three-dimensional body position and movement on a two-dimensional screen. Previous work has indicated that there is an inherent degradation in virtual spatial task speed and accuracy when completing a three-dimensional task on a standard two-dimensional display screen. This degradation has implications for the use of immersive 3D environments in movement rehabilitation, especially for elderly stroke patients who may also be suffering from cognitive as well as motor control impairments. These patients are also unlikely to have any prior experience with either virtual environments, or motion-capture type control interfaces. It is therefore desirable to find methods of reducing or removing any dimensional performance degradation caused by the interface, minimising additional patient cognitive load and maximising rehabilitation potential as much as possible.

There are innumerable design decisions to be made when representing the user's movement in a two-dimensional virtual environment, and the number of possible implementations only increase with the representation of the third dimension. The choice of movement representation and features included in the virtual environment is not inconsequential; rather, many decisions made with regard to the representation of movement and the virtual environment have the potential to affect perception of the scene and the subject's own movement. The individual's perception, in turn, affects their own movement trajectories in the task. With regard to the application of this system to a clinical rehabilitation environment, it is crucial that the key factors affecting performance first be identified. Then, these factors should be optimized to produce a virtual environment with the smallest learning curve, that is, the fewest barriers to immediate processing of the environment and the subject's own influence over the movement of the avatar.

Complicating this understanding, however, is the fact that each representational factor has the potential to influence the effects of the others on task performance. For that reason it is necessary to separate the factors of interest and manipulate them in as controlled a manner as possible. Where some potentially confounding design decisions (e.g. background colour, avatar shape) could not be avoided, the specific designs used were those deemed the simplest and least likely to influence performance or contribute to the task demands (see Methodology and Experimental Methods subsections for task design specifics). Finally, as will be evident in the more detailed description of experimental methods, the manipulation of the factors of interest took a full-factorial design where possible, meaning all combinations of levels were compared. This design provided the added benefit of testing some basic assumptions in the fields of visual perception and VR, in addition to testing for the key influential factors and VR circumstances that optimize task performance in the Z dimension. The evidence from the following experiments either for or against these assumptions will be highlighted in the results and discussion sections throughout the chapter.

The set of experiments described here addressed the separable effects of the following factors on lower-limb motor performance in a virtual environment task: mapping, grating, stereoscopic 3D, contrast, and scaling. Mapping refers to the representation of real-world X, Y and Z dimension movement as any continuous dimension on a 2D screen. The screen dimensions used for mapping the real-world movements were the two most conventional and intuitive dimensions used in the most basic VR environments: screen X (horizontal plane), and screen Y (vertical plane). The third on-screen dimensional mapping was size scaling, which mimics the relative size factor used in making near-far (i.e. Z dimension) judgments in the real world. While the goal of this experiment was to identify which type of mapping of the real-world Z dimension facilitates task performance, motion in all three real-world dimensions was mapped to all three screen representations. This allowed the verification that the real-world X and Y dimensions judged relative to the subjects' point of view are best represented by the screen X and Y dimensions. The screen-by-real-world factorial design also provided a measure of the discrepancy between performance in the real-world Z dimension compared to X and Y, that is, the main effect of real-world dimension on task performance. Performance in the real-world X and Y dimensions under the respective optimal mapping conditions can then be used as a target for performance in the Z dimension.

The second experiment presented in this chapter describes the effect of grating on task performance in each real-world dimension. The grating factor refers to the addition of background information, which may facilitate the judgment of on-screen distance, and thus the degree of real-world movement required to achieve the desired on-screen distance. While the real-world to screen dimensional mappings were tested using using a simple object representation on a solid background, it could be argued that additional background information is crucial for judging the degree of on screen movement, and thus the degree of real-world movement required to accomplish the task in the VR environment. A standard example of this is the addition of perspective lines to aid in the perception of on-screen depth. Performance in the three dimensions was tested using the three gratings that correspond to the on-screen mappings: horizontal lines (X dimension), vertical lines (Y dimension), and perspective lines (Z dimension). Each grating type was tested across multiple world-to-screen mappings in order to test the assumptions of optimal performance in the X and Y dimensions, as well as to determine whether the degradation in performance in the Z dimension is reduced with the addition of grating lines in any direction.

The third experiment tested performance on the same task with and without added stereoscopic 3D, which is a well-known and commonly-used method of representing depth on a two-dimensional surface. The illusion of depth in stereoscopic 3D is achieved using two slightly offset copies of the on-screen images, each in a contrasting colour, which are displayed to each eye when special colour filter glasses are worn. The goal of this experiment was to test the specific effect of stereoscopic 3D on task performance, and in particular to determine whether performance in the real-world Z dimension, was improved with the addition of stereoscopy depth cues. For this experiment, the mapping condition was also varied between subjects in order to identify any interactions between stereoscopic 3D and the real-world to on-screen dimensional representation. This again allowed the verification of the hypothesized performance degradation in the Z dimension relative to X and Y, and the separation of main effects of stereoscopic 3D from interactions between dimensional mapping and stereoscopy depth cues.

Yet another method of representing movement in the third dimension is change in contrast or illumination of an object. In the real world, the way that light reflects from a fixed source on to an object and onto the eye changes as the person's viewpoint or the position of objects change relative to one another. The resulting change in luminosity is one of the many cues available to the visual perception system that provides information about a relative change in position between the person and objects in his/her environment. The contrast of an object can be varied as a continuous factor, and thus it can be mapped to signal change in real-world movement. In the fourth experiment described in this chapter, the contrast of an object was mapped to the real-world movement in each dimension in order to test whether contrast in a VR task provides an added benefit to performance over a no-contrast representation, and whether the effect of contrast was more or less beneficial when mapped to any particular dimension. The key dimension of interest is again that of depth (Z), as this dimension is the most difficult to represent on a two-dimensional screen without degradation in performance. However, the contrast factor was again tested across all three dimensions in order to test the assumption that any performance benefit as a result of the additional contrast cues would be greater in the Z dimension than in X and Y.

The fifth experiment presented aimed to investigate the effect of scaling on performance in each dimension. Scaling refers to the gain factor used to adjust the subjects' movement as it is mapped to on-screen motion. As will become apparent later in this Chapter, the scaling experiment followed on the findings from the previous experiments that performance was consistently worse in the

Z dimension. This deficit is the result of misjudgments of on-screen distance in the Z dimension, of the mapping of this distance to real-world Z dimensional movement, or some combination of these two factors. After the first few experiments had confirmed the suspected performance deficit in the Z dimension, the next step was to identify the scaling parameters that correct the known error in depth distance judgment. In the fifth experiment presented here, the world-to-screen scaling factor was parametrically manipulated from range of values less than and greater than one, where one represents a direct one-to-one correspondence between movement in the real world and on-screen motion. The scaling factor was varied across multiple mapping conditions in order to test the assumption that the one-to-one scaling is optimal for representing motion in the X and Y dimensions, and that the problem of distance misjudgment in a VR environment is unique to movement in the Z dimension.

The purpose of the first five experiments was to identify the factors affecting contributing to sub-optimal performance in a lower-limb VR task, and in particular to performance in the real-world Z movement plane. The on-screen task environment for these experiments was intentionally abstract and simplified in order to separate confounding perceptual effects on visuomotor mapping and task fluency. However, performance in such tasks may not be generalisable to VR environments with more complex and realistic visual displays, where multiple perceptual cues are combined. Furthermore, as the goal of a motor rehabilitation task in a clinical setting is to generalise the practiced motions to non-clinical situations, the execution of motions with a more realistic level of complexity is likely to be beneficial for many stroke patients. Such complex tasks might involve navigation around realistic 3D objects, where there are no constraints on the dimensionality of real-world or on-screen motion. The final experiment (Section 3.7, Obstacle Avoidance) tested the application of results from previous psychophysical experiments to a more complex and realistic rehabilitation task. As the methods used in the obstacle avoidance task differ substantially from those used in the other experiments, the obstacle avoidance methods are not discussed in the following section on general methodology. Rather, these are detailed within the Obstacle Avoidance experimental section (see Section 3.7.2 for methods).

3.1 General Methodology

Overall, methods in the clinical VR field appear to be inconsistent and immature as a whole, and thus the methods in this section are presented such that they may be used as a reference framework for future researchers designing and implementing similar systems in this area. A description of and rationale for the general methods used in the psychophysical experiments will be presented in this section. The reports of each experiment will contain more specific descriptions of the methods used, and these individual experimental methods sections will detail any deviations from the general methodology.

The experiments testing the effects of mapping, grating, 3D and scaling used a "disk in hole" task. In this task, the subject must move a solid "disk", which was mapped to the subject's foot movement, into a wire-frame circle ("hole"), and maintain this position for five seconds. The targets appeared in unpredictable locations over multiple trials, and the dependent measure was the average time to completion over trials in the same condition. This is analogous to the 3-dimensional "peg in hole" task often used to assess and compare accuracy within virtual and real environments (Unger et al., 2001, 2002; Amirabdollahian et al., 2005; Burdea et al., 2000). The "peg in hole" task was adapted for use with the lower-limbs (dominant foot movement) because of the potential application of the system to gait rehabilitation in stroke patients. The decision to use lower limb movement was also made to address a gap in current research, as the majority of VR rehabilitation research is undertaken using tasks involving limbs such as the arm or hand.

Of the first five experiments presented in this chapter, one exception to the "disk in hole" method described above was that of contrast manipulation. The contrast experiment used a similar task design as that described above, with a modification to the nature of the target. In the contrast experiment, the subject was required to adjust the luminosity of the "disk" object using movement in a single dimension until it matched the luminosity of the target on that trial. The contrast version of the task was designed to be as similar as possible to the task described above, in order to allow for the comparison of results across experiments. In the version of the task used in the contrast experiment, trial completion was achieved when the target luminosity value was maintained for five seconds. Like the other "disk in hole" task described above, the dependent measure in the contrast experiment version was the speed of trial completion averaged over multiple trials.

For both the "disk in hole" task and the contrast-matching version of this task outlined above, the general methods were selected to meet two main criteria. First, the task graphics should be simple in order to test the effect of each factor separately from any others, to the extent that this is possible. Second, given that the larger aim of this research is to identify the factors that may negatively affect performance in a gait rehabilitation context, the dependent measure should index speed and/or accuracy of lower limb movement. In the experiments presented, the dependent measure was average time to trial completion (speed), although the trial completion itself required accuracy. There is likely to be some trade off between speed and accuracy in these tasks, and the balance of this trade off is reflected by the time to completion. That is, both slower and less accurate movements will result in slower times to completion. Given the accurate motion tracking abilities of the Stromohab system, it is possible to identify the nature of performance deficits by examining the full motor trajectories over trials, rather than simply the mean time to completion. However, this level of analysis is beyond the scope of the present study.

3.1.1 Coordinate Systems

Through this thesis, two primary coordinate systems are described, *World* and *Virtual*. These represent movement in the physical and virtual worlds respectively. The limits of, and transformations between these coordinate systems are shown in Figure 3.2.

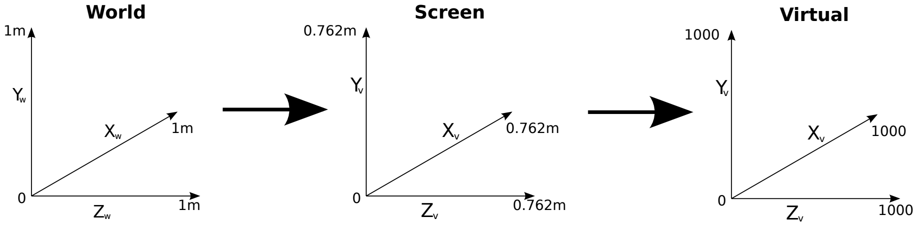

Figure 3.1: *Mapping of physical subject movement to physical distance moved on screen. Perspective vanishing point at maximum Z.*

Physical coordinates and dimensions are represented using X_w, Y_w, and Z_w notation. Similarly, coordinates and dimensions in the virtual environment are represented using X_v, Y_v, and Z_v.

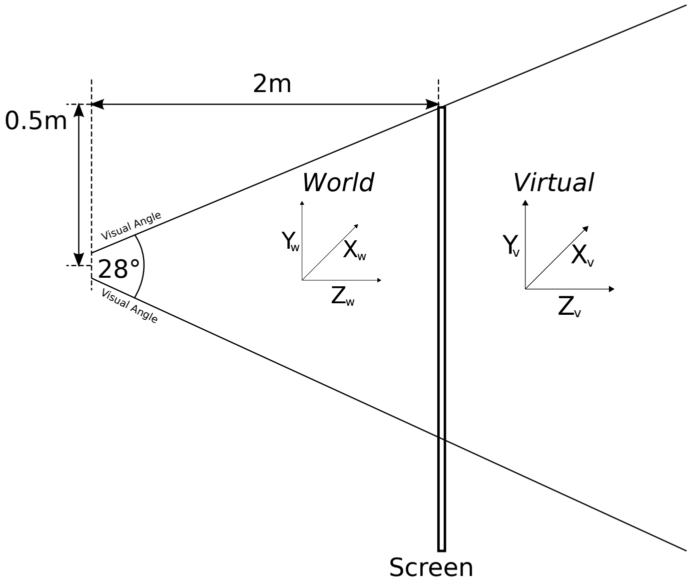

Figure 3.2: *Visual angle in experimental setup. Subject to stimulus distance=2m. Maximum subtended visual angle subtended=28°.*

3.2 Mapping

Conventionally, Z_w plane movement is represented onscreen by modifying the size of an object, emulating perspective in the real world. Judging Z_v-distance (distance “into” a scene) is difficult in virtual environments and accuracy is dependent on viewpoint, that is, the vector from users’ eyes to their avatar representation within a virtual environment (Unger et al., 2002). Fitts’ law predicts that, on a computer monitor, the time required to rapidly move to a target area is a function of the distance to the target and the size of the target (Howarth et al., 1971). This law is well tested (Stuart et al., 1978; Murata, 1999; Zhai, 2002; Accot and Zhai, 2003), and may have implications in a motion-capture driven virtual reality environment. If this is the case, the additional difficulty presented by modifying size to represent distance in a virtual environment may detract from the efficacy of some 3D rehabilitation tasks, and therefore virtual task design must be carefully considered, especially when the virtual task is being used in a rehabilitation setting.

3.2.1 Hypothesis

Within the environment, testing the accuracy of individual planes of movement will indicate whether or not there is a deficiency in the Z_v plane, as compared to the X_v and Y_v planes of movement. Furthermore, testing alternate mappings (X_w-Z_v, Z_w-Y_v, etc), will show any other more optimal mappings, should they exist.

3.2.2 Methods

Participants

In total, 16 subjects were recruited from a university cohort. The mean age was 22.77 years ($SD = 6.84$), and the group was 75% male. Participants had self-declared normal or corrected to normal vision, and none reported any vision deficiencies, including colour blindness. Subjects were not compensated for their participation.

Design

The experiment took a full factorial design. There were two within-subjects independent factors, “World” and “Screen”, each with three levels: X (left-right), Y (up-down) and Z (forwards-backwards). Each factor represents the dimension in which the subject moves (World) or experiences movement (Screen). Both “World” and “Screen” independent variables were manipulated within-subjects, meaning all subjects completed trials for all combinations of “World” and “Screen” mappings.

Procedure

Before beginning, subjects were stood 2.5m from a 50” plasma television monitor mounted vertically in front of them. All natural lighting was excluded from the room, leaving it lit by laboratory lighting only. This lighting was positioned to avoid both glare and reflections on the display screen. The height of the screen was adjusted such that the centre was at the subjects eye level, producing a visual angle of 0° within the task.

The task used was “Disk in Hole”. In this task, the subject must move a solid disk into a wire-frame circle (“hole”), and maintain this position for five seconds. A passive optical motion capture marker was attached to their dominant foot, and a solid black circular disk on a red background was displayed on screen (see Figure 3.6). Movement of the subjects dominant foot controlled the movement of this disk, transformed and constrained to a single on screen dimension as determined by the mapping condition.

Subjects initially completed a short training protocol of three targets in each of the X_w-X_v, Y_w-Y_v and Z_w-Z_v mappings to ensure they understood the task. During the task, a larger circular outline appeared sequentially in thirty predetermined locations on screen (the target), and the subject was required to move their foot and place the disk inside this outline. The order of target locations was the same for each subject. Accurate placement (within 10% of the centre of the target) was indicated by the background changing from red to green (Figure 3.7). After the subject had maintained accurate placement for 5 seconds, the target moved to a new location.

The order of mapping conditions was randomised, with each mapping condition consisting of 30 consecutive targets at unpredictable locations. All possible World-to-Screen mappings were tested; these are listed in Table 3.1. Target locations were consistent between mappings and generated prior to testing. Generation of the target locations is discussed below. To minimise fatigue effects, subjects were allowed a rest period of up to 2 minutes between each mapping presentation.

		Screen Dimension		
		X_v	Y_v	Z_v
World Dimension	X_w	$X_w X_v$	$X_w Y_v$	$X_w Z_v$
	Y_w	$Y_w X_v$	$Y_w Y_v$	$Y_w Z_v$
	Z_w	$Z_w X_v$	$Z_w Y_v$	$Z_w Z_v$

Table 3.1: *Experimental mappings*

Target Generation

To generate the unpredictable target locations for each trial, all possible target locations were generated within the virtual environment, which is a "cube" with all sides of length 1000. First, the dimensions of the cube were reduced so that the potential target locations had to fall between 200 and 800 in each dimension. These reduced dimensional limits allowed a buffer of length 200 along the edges of each dimension to avoid overlap between the edges of the stimuli and the edge of the screen. Second, all possible points were generated with a spacing of 100 between points along each dimension within the resulting cube. From these X_v, Y_v, Z_v coordinates, 30 random points were selected with the constraint that no two points should have the same consecutive dimensions. This constraint ensured that subjects never remain in the same location for more than a single target.

Statistics

Analysing the results using a two-way repeated-measures Analysis of Variance (ANOVA) will indicate the presence of main effects for either World or Screen dimensions, as well as any interaction effect. If significant effects are found, further pairwise t-tests (Bonferroni corrected for conservatism) will indicate differences between levels within each factor.

3.2.3 Results

The total completion time in each of the World-to-Screen transformation conditions was divided by the number of trials in each world-by-screen condition in order to obtain the mean completion time over trials for each subject. The subject averages were then averaged over each World-to-Screen condition, to provide a mean completion time for each World-to-Screen mapping. The means and standard deviations for all World-to-Screen mapping conditions are shown in Table 3.2.

World	Screen	Mean	Standard Deviation
X_w	X_v	9.29	1.52
X_w	Y_v	12.75	1.53
X_w	Z_v	12.42	1.39
Y_w	X_v	13.03	1.39
Y_w	Y_v	9.22	1.44
Y_w	Z_v	9.85	1.4
Z_w	X_v	10.76	1.42
Z_w	Y_v	9.89	1.58
Z_w	Z_v	11.31	1.39

Table 3.2: *Means and standard deviations of single-trial completion times (in seconds) for each World-to-Screen dimensional mapping condition*

The data were tested and found to have met the assumptions of Mauchly's test of sphericity. The tests of sphericity for both within-subject factors and the interaction were non-significant (see Table 3.3, meaning the data has not violated the assumptions for a two-way repeated-measures ANOVA.

Within-Subject Effect	Mauchly's	Degrees of Freedom	Significance
World	0.979	2	.862
Screen	0.807	2	.222
World-by-Screen	0.478	9	.363

Table 3.3: *Results of Mauchly's test of sphericity for the two within-subjects factors and the interaction*

ANOVA Results

A two-way repeated-measures ANOVA was then conducted to evaluate the effect of transforming real-world movement to on-screen movement on mean target completion time, and to isolate the main effects of both World and Screen planar movements. Both the World and Screen main effects were significant, World dimension: $F(2, 30) = 95.40, p < .001$, Screen dimension: $F(2, 30) = 26.29, p < .001$. There was also a significant interaction effect present, World-by-Screen: $F(4, 60) = 442.14, p < .001$.

World Main Effect

The significant main effect of World was followed up with pairwise comparisons to determine the significant differences between levels of the World factor. The alpha significance level for these tests was Bonferroni-corrected to $p < .017$. As can be seen from Figure 3.3, when moving in the real-world X_w dimension, the mean target completion time was significantly longer than both Y_w ($X_w - Y_w$ mean difference: 0.79, standard error: 0.07, $p < .001$), and Z_w ($X_w - Z_w$ mean difference: 0.84, standard error: 0.06, $p < .001$). The mean difference between Y_w and Z_w real-world movement was not significant.

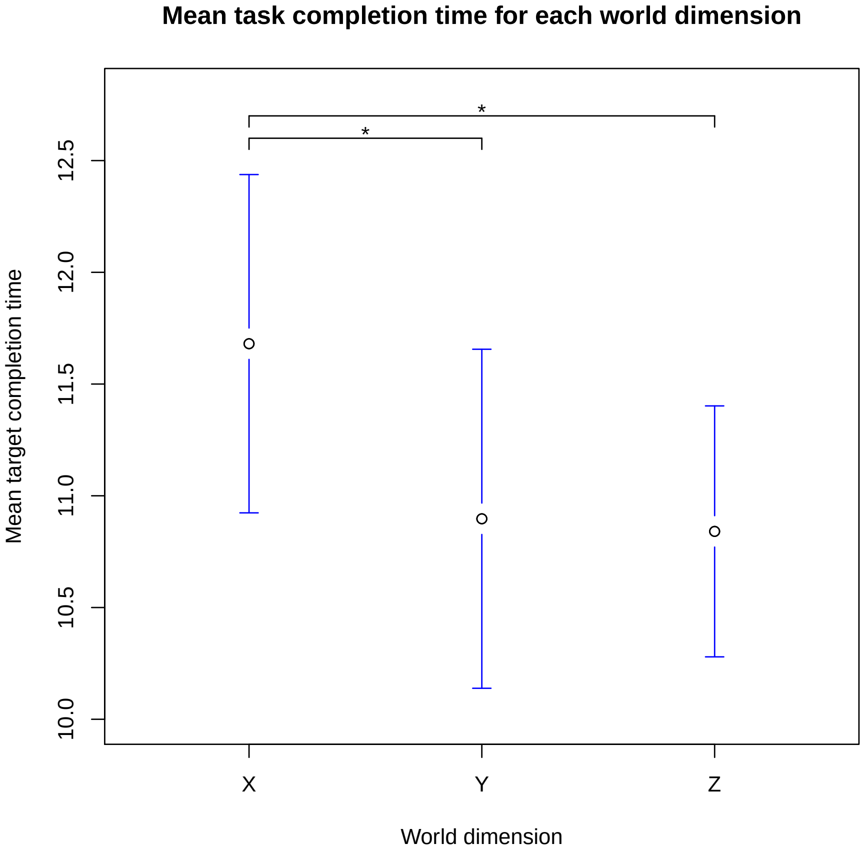

Figure 3.3: *Main effect of real-world dimension on mean trial completion times (in seconds). Error bars represent 95% confidence intervals. " * " indicates a significant difference between the conditions.*

Screen Main Effect

The significant main effect of Screen was followed up with pairwise comparisons to determine the significant differences between levels of the Screen factor. The alpha significance level for these tests was Bonferroni-corrected to .017. Figure 3.4 shows that trials mapped to both X_v and Z_v screen dimensions took significantly more time to complete than trials mapped to the Y_v screen dimension ($X_v - Y_v$ mean difference: 0.40, standard error: 0.10, $p < .001$, $Z_v - Y_v$ mean difference: 0.57, standard error: 0.07, $p < .001$). The mean difference between X_v and Z_v screen mappings did not reach the Bonferroni-corrected significance threshold of .017.

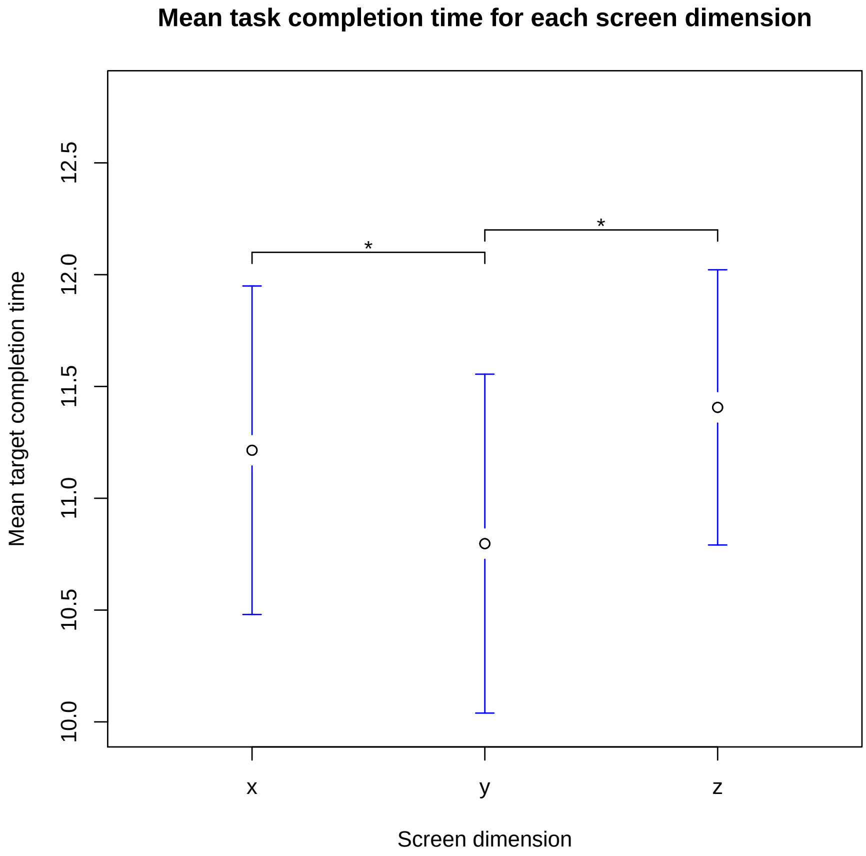

Figure 3.4: *Main effect of screen dimension on mean trial completion times (in seconds). Error bars represent 95% confidence intervals. " * " indicates a significant difference between the conditions.*

World-by-Screen Interaction

The World-by-Screen interaction plot in Figure 3.5 shows the significant interaction effect found in the two-way repeated-measures ANOVA described above. Further investigation into this interaction is presented below.

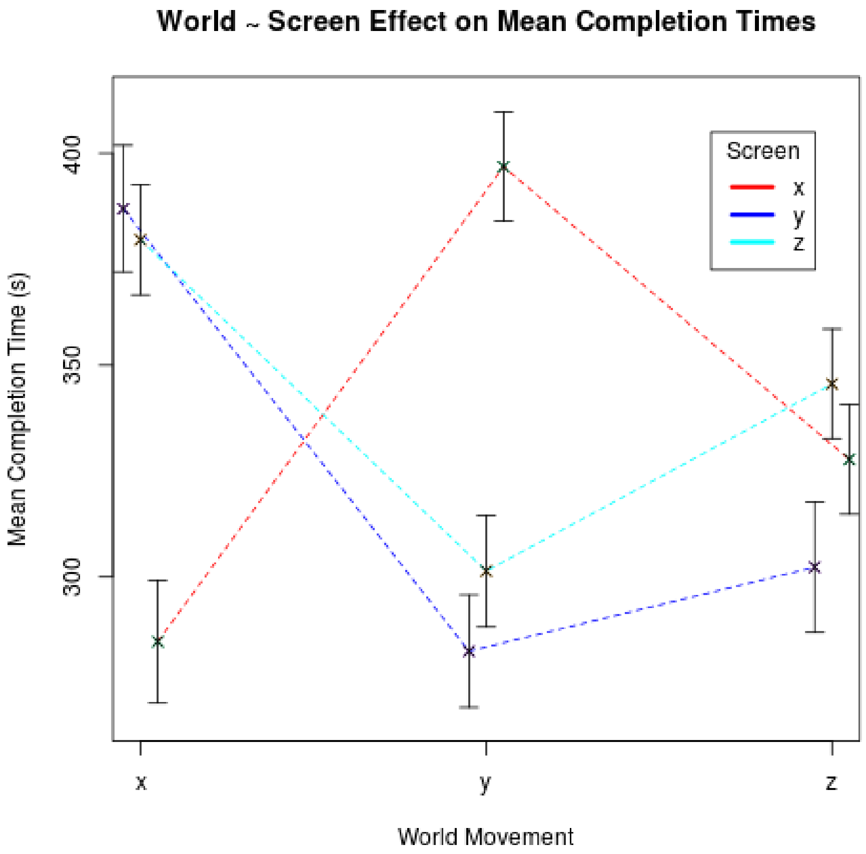

Figure 3.5: *Effect of interactions between of real-world and screen dimensions on total condition completion times (in seconds). Error bars represent 95% confidence intervals.*

Pairwise t-tests were performed within the X_w condition with a Bonferroni-adjusted alpha of $p < .0057$ to correct for mulitple comparisons. The results showed that mapping X_wX_v was the most optimal mapping for X_w, with a significantly lower mean target completion time than X_wY_v and X_wZ_v. There was no significant difference between X_wY_v and X_wZ_v.

Pairwise t-tests were performed within the Y_w condition with a Bonferroni-adjusted alpha of $p < .0057$ to correct for mulitple comparisons. The results showed that there were significant differences between each screen level. $Y_w Y_v$ was the most optimal mapping, followed by $Y_w Z_v$, then $Y_w X_v$.

Pairwise t-tests were performed within the Z_w condition with a Bonferroni-adjusted alpha of $p < .0057$ to correct for mulitple comparisons. The results showed that there were significant differences in mean trial completion time between each screen level. $Z_w Y_v$ was the most optimal mapping, followed by $Z_w X_v$, then $Z_w Z_v$.

"Straight" Mappings

The mean trial completion times for "straight" mappings (i.e. World dimension mapped to same Screen dimension) were compared to test the hypothesis that there was a performance deficiency in the Z_v plane compared to the X_v and Y_v planes. Pairwise t-tests with alpha adjusted to $p < .0057$ showed that mean completion time in $Z_w Z_v$ was significantly greater than in $X_w X_v$ and $Y_w Y_v$. $X_w X_v$ and $Y_w Y_v$ mean trial completion times were not significantly different.

Optimal Mappings

For each World dimension, the mean trial completion times for optimal mappings (i.e. mapping to the Screen dimension where performance was best) were compared to test whether a performance deficiency in any World dimension exists when it is represented with the optimal Screen mapping. The optimal mapping for the X_w dimension was X_v ($X_w X_v$), and the optimal mapping for the Y_w dimension was Y_v ($Y_w Y_v$). For movement in the Z_w dimension, the optimal mapping was Y_v ($Z_w Y_v$). Pairwise t-tests with alpha adjusted to $p < .0057$ showed that the mean trial completion time in $Z_w Y_v$ was significantly greater than $Y_w Y_v$. There was no significant difference in mean completion times between $X_w X_v$ and $Y_w Y_v$, or between $X_w X_v$ and $Z_w Y_v$.

Mapping Results Summary

The results of the present experiment showed that both World and Screen dimensional mappings had a significant effect on mean trial completion times, and the interaction between these two factors was also significant. Contrasts between the levels of World revealed that mean trial completion times were significantly higher when subjects were moving in the X_w plane (left-right) than while moving in either the Y_w (up-down) and Z (forward-backward) planes. This means that, irrespective of the Screen condition, performance was generally worse for real-world movement in the X_w dimension. Contrasts within the Screen factor showed that, irrespective of World dimension, the mean trial completion times were faster when the subjects' motion was mapped to the Y_v dimension than to either the X_v or Z_v dimensions.

Although the main effects of both factors were significant, the effects of World and Screen dimensional mapping are best understood in terms of the interaction between the two factors. In both the X_w and Y World dimensions, performance was best when the World dimension was mapped to its corresponding Screen dimension, i.e. X_w to X_v and Y_w to Y_v. This was not the case in the Z_w condition, where performance was best when Z_w was mapped to Y_v. The optimal mappings are summarised in Table 3.4. The z_w to z_v mapping was found to be suboptimal, and thus not shown.

World	Screen
X_w	X_v
Y_w	Y_v
Z_w	Y_v

Table 3.4: *Optimal mappings*

3.3 Grating

The results of the mapping experiment (Section 3.2) confirmed the hypothesized performance deficit for the Z_v (depth) dimension relative to X_v and Y_v. However, in the the mapping experiment, the only objects displayed on the screen were the user-controlled object ("disk") and a target location ("hole"). In real-world perception, the use of object size to judge the distance of the object in the Z_w-dimension is determined relative to the surrounding static environmental cues. It could be argued that the presence of additional visual information for reference provides a specific performance benefit for the on-screen Z_v dimensional mapping, where the subject must map changes in object size to movement. To test whether performance is improved with the addition of visual "reference" information, linear gratings were added to the background in the task. There were three grating conditions corresponding to the directions of movement for each mapping, as well as one no-grating ("None") control condition.

The results of the mapping experiment (Section 3.2.3) also showed that performance was best for "straight mappings" (dimension mapped to self) in X_w

and Y_w, and that performance in the Z_w dimension was optimized when it was mapped to the Y_v dimension on screen. Rather than testing the full set of 9 World-by-Screen mapping combinations, in this experiment the levels of the mapping factor were limited to X_wX_v, Y_wY_v, Z_wZ_v and Z_wY_v. The inclusion of the X_wX_v and Y_wY_v optimal mappings allowed for comparisons between performance in these dimensions and in Z_v, as one of the aims of the current research is to reduce or eliminate the depth-specific performance deficit in VR environments. Both Z_wY_v and Z_wZ_v mappings were included in the present experiment in order to confirm the finding from the mapping experiment that performance in the Z_w dimension is improved when mapped to the Y_v dimension. It was not clear from the results of the single mapping experiment that the superiority of Z_wY_v performance over Z_wZ_v would persist with the addition of background gratings for visual reference.

3.3.1 Hypothesis

It is hypothesized that increasing the amount of on-screen spatial information will improve performance in the task. Specifically, gratings perpendicular to the direction of movement will provide additional motion information, and therefore comparison of horizontal, vertical and diagonal background gratings with a no-grating control condition will show an improvement in an associated dimension, i.e. a vertical grating will improve performance in the X_wX_v mapping, a horizontal grating will improve performance in the Y_wY_v mapping, and a diagonal grating matched to the perspective of the Z_wZ_v mapping will improve performance in that mapping. As the previous task showed that a Z_wY_v mapping lead to better task performance than Z_wZ_v, the Z_wY_v mapping will also be tested. It is hypothesized that, as in the mapping experiment, the mean trial completion times will be shorter in Z_wY_v than in Z_wZ_v. It is further suggested that, as for Y_wY_v, a horizontal grating will improve performance in the Z_wY_v mapping.

3.3.2 Methods

Participants

In total, 20 subjects were recruited from a university cohort. The mean age was 24.42 years ($SD = 7.69$), and the group was 95% male. Participants had normal or corrected to normal vision, and none reported any vision deficiencies, including colour blindness. Subjects were not compensated for their participation.

Design

The mapping factor was varied between-subjects, with each subject randomly assigned to one of four mapping groups: 'X_wX_v', 'Y_wY_v', 'Z_wZ_v' and 'Z_wY_v'. Within each mapping group, all four grating levels were completed: 'None', 'Horizontal', 'Vertical' and 'Perspective'. In order to control for learning or other order effects, grating presentation order was randomised for each subject. The dependent measure was mean time to target completion, defined as the total time taken to complete all targets within a grating condition divided by the number of presented targets.

Procedure

Subjects were stood 2.5m from a 50" plasma television monitor mounted vertically in front of them. A passive optical motion capture marker was attached to their dominant foot, and a solid black circular disk on a red background was displayed on screen (see Figure 3.6). Movement of the subject's dominant foot controlled the movement of this disk, constrained to a single on screen dimension as determined by the mapping condition.

During the task, a larger circular outline (the target) appeared sequentially in 30 predetermined locations on screen, and the subject was required to move his/her foot in order to place the solid black circle (the disk) inside this outline. The order of target locations was the same for each subject. Accurate placement (within 10% of the centre of the target) was indicated by a red to green change in background colour (Figure 3.7). After the subject had maintained accurate placement for 5 seconds, the target moved to a new location. After completion of a grating condition, subjects were allowed up to a five minute break before continuing. Subjects completed a demonstration target in each grating condition prior to beginning the experiment. Sample grating images can be found in Figure 3.8.

Statistics

Analysing the results using a two-way mixed design ANOVA will indicate the presence of main effects for either mapping or grating factors, as well as any

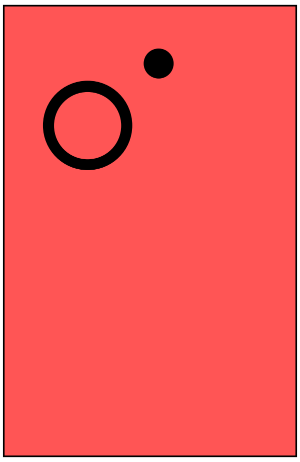

Figure 3.6: *Screen view with disk outside target (not to scale)*

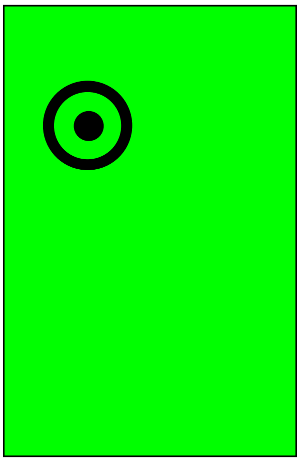

Figure 3.7: *Screen view with disk inside target (not to scale)*

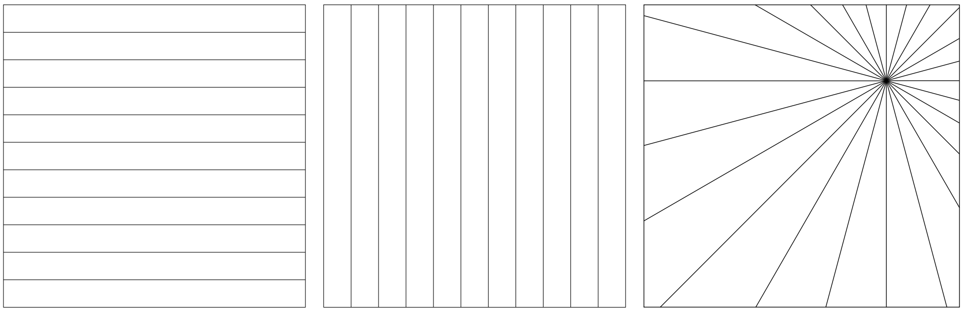

(a) *Horizontal* **(b)** *Vertical* **(c)** *Perspective*

Figure 3.8: *Example Gratings (not to scale)*

interaction effect. If significant effects are found, further pairwise t-tests (Bonferroni corrected for conservatism) will indicate differences between levels within each factor.

3.3.3 Results

For each subject, the total completion times within each of the four grating conditions were divided by the total number of trials in each condition in order to obtain the mean completion times per trial in each grating condition. These subject mean completion times were then averaged across the five subjects in each of the mapping conditions to obtain a grand average single-trial completion time for each mapping-by-grating condition. These averages and standard deviations can be found in Table 3.5.

Mapping	Grating	Mean	Standard Deviation
X_wX_v	None	7.82	0.34
	Horizontal	7.11	0.32
	Vertical	6.64	0.35
	Perspective	7.42	0.35
Y_wY_v	None	7.93	0.69
	Horizontal	6.51	0.33
	Vertical	7.24	0.45
	Perspective	7.78	0.43
Z_wZ_v	None	9.60	0.68
	Horizontal	9.47	0.60
	Vertical	9.41	0.23
	Perspective	8.93	0.18
Z_wY_v	None	8.19	0.54
	Horizontal	7.41	0.42
	Vertical	8.32	0.29
	Perspective	7.70	0.54

Table 3.5: *Means and standard deviations of single-target completion times (in seconds) for each mapping-by-grating condition*

The data were tested and found to have met the assumptions of Mauchly's test of sphericity for the within-subjects independent variable (grating), $p = .057$, and of Levene's test of homogeneity of variance for the between-subjects independent variable (mapping, Table 3.6). All levels of the within-subjects factor are shown in Table 3.6 as the homogeneity of variance across the levels of the mapping factor were tested within all levels of the grating factor. The non-significant p-values indicate that the distributions of variance across all levels of mapping within all levels of grating do not differ from a normal distribution.

ANOVA Results

A two-way mixed-design ANOVA was conducted to evaluate the overall effects of dimensional World-to-Screen mapping and the background screen gratings lines on the mean trial completion times. All effects are reported as significant at $p < .05$. The ANOVA results showed that there was a significant main effect of grating on mean target completion time, $F(3, 48) = 10.32$, $p < .001$, and a significant main effect of mapping on mean target completion time, $F(3, 16) = 90.64$, $p < .001$. There was also a significant interaction between grating and mapping factors, $F(9, 48) = 4.45$, $p < .001$.

Main Effect of Grating

The significant main effect of the grating factor was followed up with planned contrasts to determine which levels of grating were significantly different from the no-grating control condition ("None"). Planned contrasts revealed that, irrespective of the mapping condition, the Horizontal, $F(1, 16) = 18.05$, $p = .001$, Vertical, $F(1, 16) = 21.90$, $p < .001$ and Perspective, $F(1, 3) = 7.16$, $p = .017$ gratings all improved the target completion time relative to the control grating condition, (Figure 3.9).

Main Effect of Mapping

The significant main effect of mapping was followed up with contrasts to determine which levels of mapping were significantly different from one another. Contrasts revealed that performance in the Y_wY_v mapping was significantly better than the mean performance across all mappings, $p < .001$, and that performance in the Z_wZ_v mapping was significantly worse than the mean, $p < .001$. Further to this, as shown in Figure 3.10, performance in the X_wX_v mapping was better than Z_wY_v, $p < .001$ and Z_wZ_v, $p < .001$. Performance in Y_wY_v was also better than Z_wY_v, $p = .002$ and Z_wZ_v, $p < .001$. Finally, performance in Z_wY_v

Grating Condition	F	df1	df2	Significance
None	1.644	3	16	.219
Horizontal	1.577	3	16	.234
Vertical	1.511	3	16	.250
Perspective	2.036	3	16	.149

Table 3.6: *Levene's test of equality of variances for the between-subjects mapping factor*

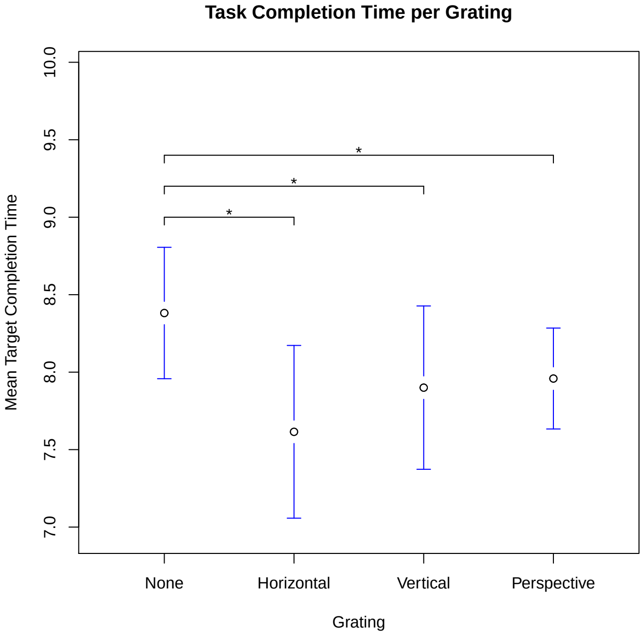

Figure 3.9: *Mean time to completion (in seconds) for each on-screen grating condition. ‘ * ’ indicates significance at $p < .05$. Error bars represent 95% confidence intervals.*

was also better than $Z_w Z_v$, $p < .001$, and there was no significant difference in performance between the $X_w X_v$ and $Y_w Y_v$ mapping conditions.

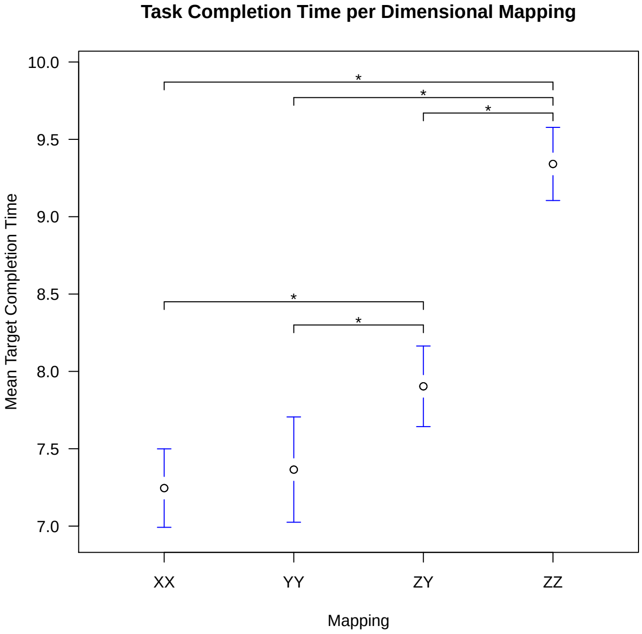

Figure 3.10: *Mean time to completion (in seconds) for each world-to-screen mapping. ' * ' indicates significance at* $p < .05$*. Error bars represent 95% confidence intervals.*

Mapping-by-Grating Interaction

The interaction effect between grating and mapping was significant, which indicates that the effect of the on-screen grating differed across World-to-Screen mapping conditions. To break down this interaction, contrasts were performed using a Bonferroni-corrected significance threshold of .0125. Each grating type was compared to the control grating ("None") across all mappings. Figure 3.11 shows the mean target completion time of each grating across each mapping group.

Contrasting the levels of gratings within each mapping shows that the Vertical grating significantly improved performance in the $X_w X_v$ mapping, $p < .001$ (Figure 3.12a), and the Horizontal grating significantly improved performance in the $Y_w Y_v$ mapping, $p = .001$ (Figure 3.12b). In the $Z_w Z_v$ mapping, no grating had a significant effect on performance (Figure 3.12c). In the $Z_w Y_v$ mapping, the difference between the control and Horizontal grating approached

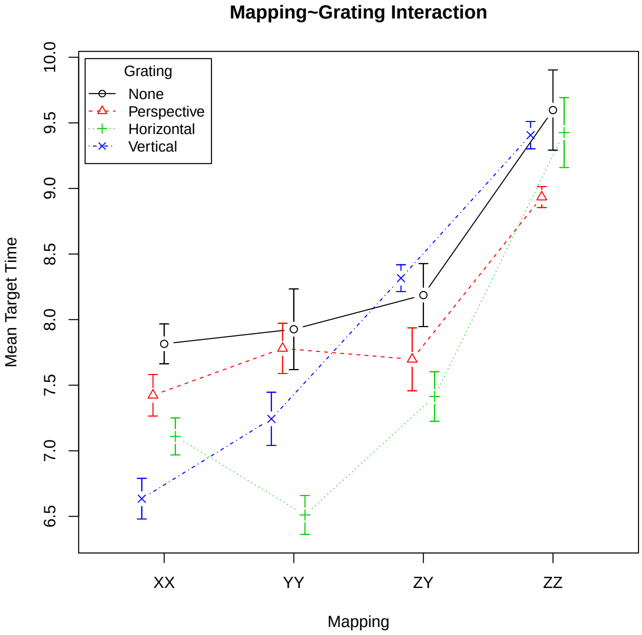

Figure 3.11: *Interaction effects between World-to-Screen dimensional mapping and on-screen grating on mean trial completion times (in seconds). Error bars represent 95% confidence intervals.*

significance ($p = .048$) but did not reach the Bonferroni-corrected .0125 alpha level (Figure 3.12d).

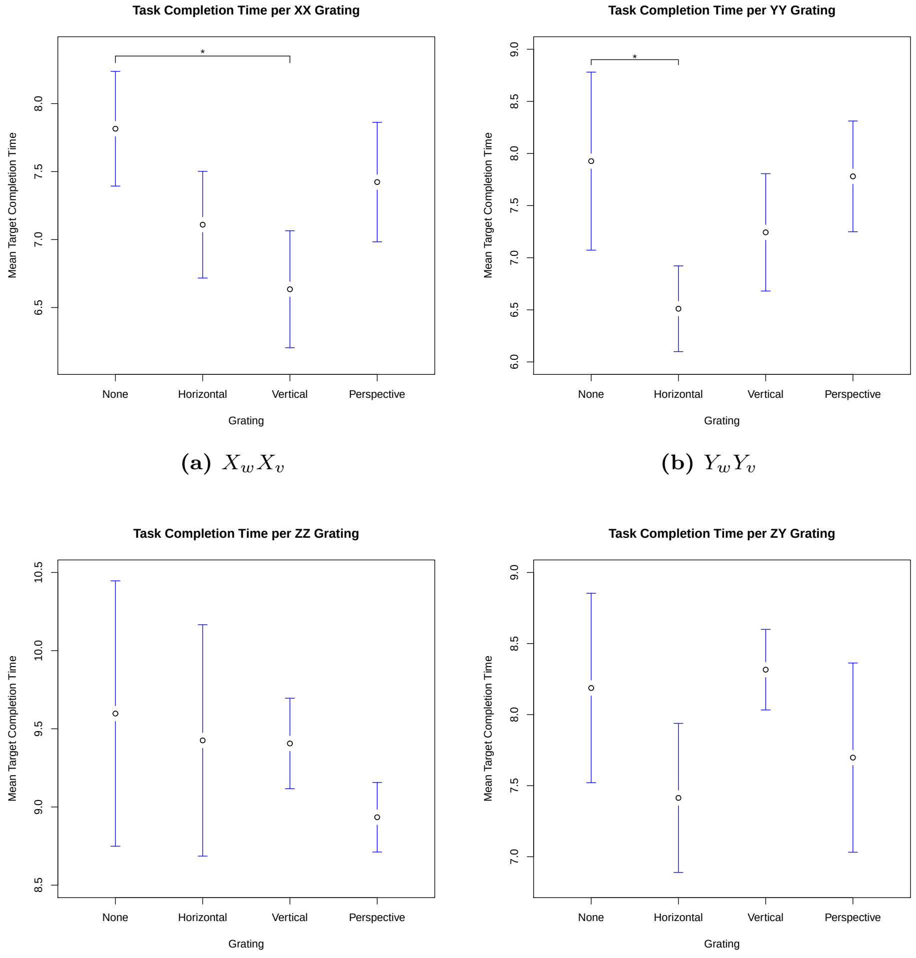

(a) X_wX_v **(b)** Y_wY_v

(c) Z_wZ_v **(d)** Z_wY_v

Figure 3.12: *Significant differences in mean trial completion times (in seconds) within each mapping-by-grating condition. In Figure 3.12d, 'Horizontal'<'None' approached significance ($p = .048$) but did not reach the Bonferroni-adjusted alpha of .0125. Error bars represent 95% confidence intervals. ' * ' indicates significance at $p < .05$.*

Results Summary

The results from the present experiment showed significant main effects of both mapping and grating on mean trial completion times, as well as a significant interaction between these two factors. Irrespective of the mapping condition, the

Horizontal, Vertical and Perspective grating conditions were found to improve performance relative to the no-grating condition.

Consistent with the results of the prior mapping experiment, performance in the Z_w to Y_v mapping condition was found to be superior to that in Z_w to Z_v. While performance in Z_w was improved when mapped to Y_v, there was still an overall degradation in performance when subjects were moving in the optimal Z_w dimension (Z_wY_v) relative to X_w (X_wX_v) and Y_w (Y_wY_v) movement.

The comparisons between mapping-by-grating conditions revealed that, as predicted, the performance was best within each mapping condition when the background grating was perpendicular to the direction of on-screen movement. The mean trial completion time was lowest in X_wX_v with Vertical gratings, in Y_wY_v with Horizontal gratings, and in Z_wY_v with Horizontal gratings. Contrary to expectations, performance in Z_wZ_v was not significantly improved with the addition of Perspective gratings.

3.4 Stereoscopic 3D

The widespread availability of 3D display products has been steadily rising in recent years, and can reasonably be expected to continue to do so. Despite the proliferation of modern 3D technology, the stereo technique for enhancing the illusion of depth and protrusion on a two-dimensional surface has remained the same since the first stereo drawings made by the sixteenth century Florentine painter Jacopo Chimenti; principally the technique of displaying a different image to each eye. This technique is called stereoscopic 3D. One common implementation of stereoscopy is anaglyph colour 3D, where the three-dimensional effect is achieved by using different colour filters over the left and right eyes in order to separately encode the associated offset two-dimensional images. When the viewer wears anaglyph glasses, each of the two offset images reaches only one eye, producing the illusion of the binocular disparity that occurs when viewing objects in the real world.

Research into the effect of stereoscopic viewing on motor performace in VR environments has produced mixed results at best (Armbrüster et al., 2008; Sandeep and Mindy; Ustinova et al., 2011; Steinisch et al., 2012), but the specific, seperable effect of 3D viewing on perceptual and motor accuracy are still not yet well understood. Therefore, the use of anaglyph 3D in a motor rehabilitation task has the potential to improve performance in the Z_v dimension, and it has the advantage of being relatively easy to add to an existing task that uses a standard two-dimensional display. Anaglyph 3D can be achieved by creating two copies of the stimuli, slightly offset and re-coloured, and requiring the viewer to wear a pair of colour anaglyph glasses. However, it is unclear as to whether the addition of anaglyph 3D provides an advantage for Z_v lower-limb motor accuracy in a VR task. Even if such a Z_v advantage exists, it is possible that motor accuracy in the X_v and Y_v dimensions suffers with the addition of anaglyph 3D, and that any performance advantage in the Z_v dimension is offset by performance degradations in X_v and Y_v.

The purpose of this experiment was to test the effects of anaglyph 3D on task performance, as measured by mean trial completion times. Specifically, the aim of the experiment was to determine whether the addition of anaglyph 3D improved motor accuracy in Z_w dimensional movement, and whether performance in the X_w and Y_w dimensions would differ with the addition of anaglyph 3D. The effects of 3D were tested across the world-to-screen mapping conditions X_wX_v, Y_wY_v, Z_wZ_v, and Z_wY_v. Previous experiments have shown that performance in the Z_w dimension is optimal when mapped to the Y_v dimension (see Section 3.2.3 and Section 3.3.3). Therefore, a secondary purpose of the experiment was to replicate this finding.

3.4.1 Hypothesis

Primary

If anaglyph 3D adds useful depth information to a scene, mean target completion time will decrease in the Z_wZ_v mapping condition. Furthermore, should enough useful depth information be added by anaglyph 3D cues, mean target completion time in the Z_wZ_v mapping will not be significantly different than mean target completion time in X_wX_v and Y_wY_v mapping conditions.

Secondary

The No-3D condition is essentially a control condition, where subjects complete the same task as that used in the previous experiments. Results in the No-3D condition are therefore expected to be consistent with other experiments utilising mapping as a variable (as described in Section 3.2.3 and Section 3.3.3). Specifically, in the No-3D condition, it is expected that the mean trial completion time in the Z_wZ_v mapping will be significantly greater than those in X_wX_v, Y_wY_v and Z_wY_v mapping conditions.

3.4.2 Methods

Participants

In total, 28 subjects were recruited from a university cohort. The mean age was 26.48 years ($SD = 4.93$), and the group was 86% male. Participants had self-declared normal or corrected to normal vision, and none were aware of any vision deficiencies. This was confirmed using a visual acuity test, in which all subjects were assessed to have a VA of at least 20/20 at 3 meters. Additionally, all participants completed the McGill University 3D vision test, and were assessed to have stereo acuity within the normal range of 10 and 50 arc seconds at 32 inches.

Design

A mixed design was used, with a between-subjects factor of mapping, and a within-subjects factor of 3D condition. Each subject was randomly assigned to one of four mapping groups: 'X_wX_v', 'Y_wY_v', 'Z_wZ_v', or 'Z_wY_v'. Within each mapping group, subjects completed both the No-3D and With-3D conditions. As in the previous experiments, the dependent variable was mean time to target completion, defined as the time taken to complete all targets within a mapping-by-3D condition divided by the number of presented targets. To minimise any order effects, the presentation order of the two 3D conditions was chosen at random for each subject.

Procedure

The basic task remains as described in Section 3.2.2, but with a change to the background colours to prevent interference with the red-blue anaglyph colours used for the left and right eyes respectively. To indicate that the user-controlled circle ("disk") was incorrectly positioned, the background was coloured purple, and to indicate that the disk was correctly positioned, the background colour changed to yellow. These colours were chosen as they are orthogonal to red and blue on a colour wheel, and thus easily distinguishable through red-blue 3D glasses.

Subjects initially completed a short training protocol of three targets. During training, the mapping condition was consistent with the assigned experimental condition, and the 3D condition was selected to be the second of the assigned 3D presentation order, to minimise the effect of learning on the experimental outcome. Within each mapping condition, the main presentation consisted of 30 targets for both 3D and No-3D conditions. Target locations were consistent with previous experiments, and were generated as described in Section 3.2.2. To minimise fatigue, there was a break of at least three minutes between the training and experimental presentations, and subjects were allowed a rest period of up to five minutes between each 3D condition.

Statistics

It is expected that a two-way mixed design ANOVA will show a main effect of mapping, consistent with prior experiments. Contrasts between the four mapping levels are also expected to maintain consistency with prior experiments. The two-way ANOVA will indicate whether there exists a main effect of 3D, should there be a significant difference between the 3D and No-3D conditions. Finally, this ANOVA will demonstrate any interaction effect between mapping and 3D factors. Should all hypothesis hold true, it is expected that there will be significant effects in all three of the above effects. Any significant effects will then be further investigated using post-hoc comparisons, with Bonferroni adjusted significance levels where required.

3.4.3 Results

For each subject, the total completion times within the No-3D and 3D conditions were divided by the total number of trials to obtain the mean trial completion times in both the 3D and No-3D conditions. These subject mean trial completion times were then averaged across the seven subjects in each of the mapping conditions to obtain a grand mean trial completion time for each mapping-by-3D condition. The resulting means and standard deviations can be found in Table 3.7.

The data were tested and found to have met the assumptions of Levene's test of homogeneity of variance for the between-subjects independent variable (mapping, see Table 3.8). Both levels of the within-subjects factor are shown in Table 3.8 as the homogeneity of variance across the levels of the mapping factor were tested within both levels of the 3D factor. The non-significant p-values indicate that the distributions of variance across all levels of mapping within both levels of 3D do not differ from a normal distribution.

ANOVA Results

Results of a two-way mixed ANOVA showed that there was no main effect of the 3D factor, $F(1, 24) = 0.34$, $p = .568$, and that there was no mapping-by-3D interaction effect, $F(3, 24) = 0.12$, $p = .946$ (see Figure 3.13).

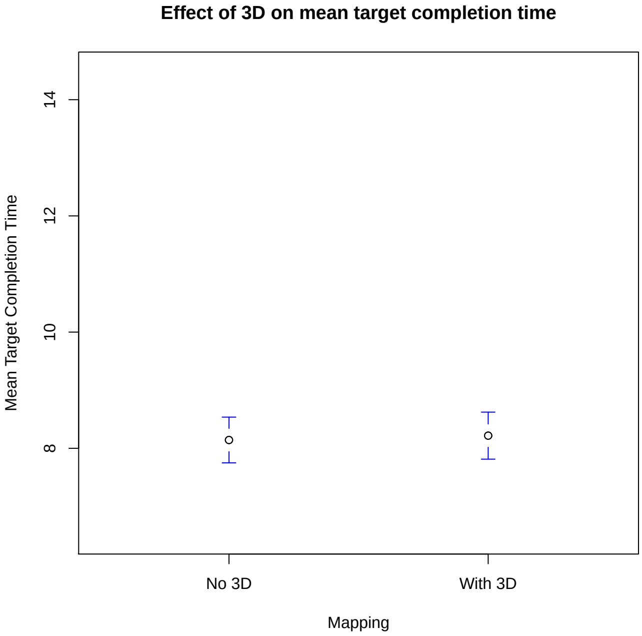

Figure 3.13: *Main effect of 3D on mean trial completion times (in seconds). Error bars represent 95% confidence intervals.*

There was a significant main effect of the between-subject mapping factor, $F(3, 24) = 89.77$, $p < .001$ (see Figure 3.14). This result was followed up

Condition	Mapping	Mean	Standard Deviation
No 3D	$X_w X_v$	7.55	0.42
	$Y_w Y_v$	7.22	0.41
	$Z_w Y_v$	8.14	0.33
	$Z_w Z_v$	9.65	0.35
3D	$X_w X_v$	7.77	0.47
	$Y_w Y_v$	7.24	0.38
	$Z_w Y_v$	8.21	0.60
	$Z_w Z_v$	9.67	0.58

Table 3.7: *Means and standard deviations of single-trial completion times (in seconds) for all mapping-by-3D conditions*

with pairwise comparisons in order to identify which levels of mapping were significantly different.

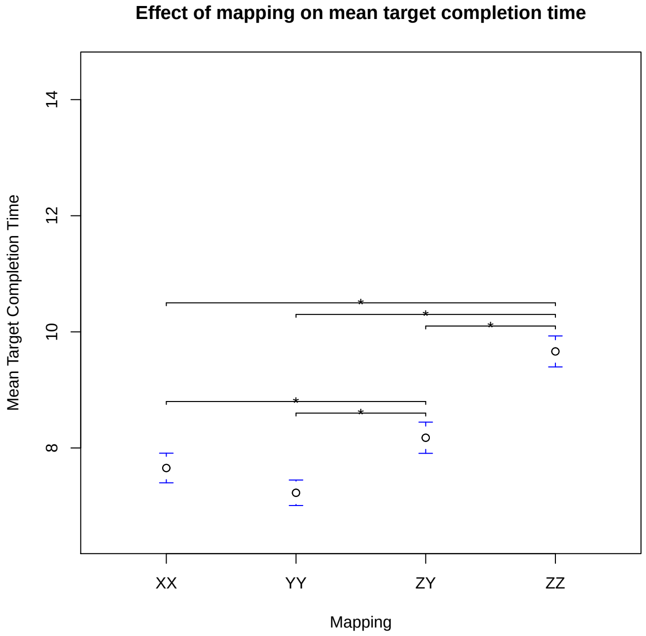

Figure 3.14: *Main effect of mapping on mean target completion times (in seconds). Error bars represent 95% confidence intervals. ' * ' indicates significance at* $p < .05$.

Mapping Pairwise Comparisons

Table 3.9 shows the pairwise comparisons between mappings for both 3D and No-3D conditions. Whilst there was no significant main effect of 3D, contrasts between mappings were consistent with previous experiments (see Figure 3.14). Mean completion time in Z_wZ_v was significantly greater than those in all other mappings ($p < .001$ for Z_wZ_v compared to X_wX_v, Y_wY_v, and Z_wY_v). Also consistent with the results of prior experiments, the mean trial completion time in Z_wY_v was significantly greater than those in X_wX_v ($p = .007$) and Y_wY_v ($p < .001$).

Condition	F	df1	df2	Significance
No 3D	.024	3	24	.995
3D	1.131	3	24	.356

Table 3.8: *Levene's test of equality of variances for the between-subjects mapping factor*

3D Condition	Mapping (I)	Mapping (J)	Mean Difference (I-J)	Significance
No-3D	$X_w X_v$	$Y_w Y_v$	0.33	.116
		$Z_w Y_v$	-0.59	.007
		$Z_w Z_v$	-2.10	.000
	$Y_w Y_v$	$X_w X_v$	-0.33	.116
		$Z_w Y_v$	-0.92	.000
		$Z_w Z_v$	-2.43	.000
	$Z_w Y_v$	$X_w X_v$	0.59	.007
		$Y_w Y_v$	0.92	.000
		$Z_w Z_v$	-1.51	.000
	$Z_w Z_v$	$X_w X_v$	2.10	.000
		$Y_w Y_v$	2.43	.000
		$Z_w Y_v$	1.51	.000
With-3D	$X_w X_v$	$Y_w Y_v$	0.52	.070
		$Z_w Y_v$	-0.45	.116
		$Z_w Z_v$	-1.91	.000
	$Y_w Y_v$	$X_w X_v$	-0.52	.070
		$Z_w Y_v$	-0.97	.002
		$Z_w Z_v$	-2.44	.000
	$Z_w Y_v$	$X_w X_v$	0.45	.116
		$Y_w Y_v$	0.97	.002
		$Z_w Z_v$	-1.46	.000
	$Z_w Z_v$	$X_w X_v$	1.91	.000
		$Y_w Y_v$	2.44	.000
		$Z_w Y_v$	1.46	.000

Table 3.9: *Mapping-by-3D pairwise comparisons showing differences in mean trial completion times (in seconds). Significance values shown are Bonferroni-corrected and can therefore be interpreted as significant at* $p < 0.05$.

Results Summary

The addition of anaglyph 3D did not have any effect on mean target completion time, and there was no difference in the effect of anaglyph 3D across the different World-to-Screen dimensional mappings. The significant main effect of World-to-Screen mapping was consistent with previous experiments. Specifically, mean trial completion time is significantly longer in the Z_wZ_v mapping compared to all other mappings, and the mean trial completion time in the Z_wY_v mapping was significantly longer than the X_wX_v and Y_wY_v mappings. Performance in the Z_w dimension was again shown to be significantly faster when this movement was mapped to Y_v as compared to Z_v.

3.5 Contrast

Light, and specifically the physics of light, is a complex subject, and there are a number of computer graphics approaches to simulating lighting within a three dimensional environment or scene. Most of these approaches use shading, or illumination models, which provide approximations of the physical laws governing light. These vary in their levels of realism and complexity. At a high level, illumination models can be viewed as generating the colour of an object's surface, at every given point on that surface.

The simplest shading model is a simple constant illumination. This assumes objects are self illuminating, and hence have their colours defined by the graphic designer. Starting from the principle that:

$$I = value \tag{3.1}$$

Where I represents the illumination colour intensity and $value$ represents the expression resulting in that value, we can define the simple constant illumination model as:

$$I = k_i \tag{3.2}$$

where k represents the basic object intensity. While this simple model is not useful in many three dimensional VR applications, it does allow the colouring of graphical objects, and hence provides the minimum amount of illumination information required to investigate scene contrast as a spatial accuracy cue within the virtual environment.

Colour

There are many computational colour models used to model and describe colour. The Hue, Saturation, and Brightness (HSB) model provides a device independent way to describe colour, and can be visualised as an upside-down cone (Figure 3.15). The three components of the HSB model are:

- **Hue**: The colour, measured in angular degrees counter-clockwise. It starts and ends at red = 0° or 360° (with yellow at 60°, green at 120°, etc.).
- **Saturation**: The purity of the colour, measured in percent from the centre of the cone (0) to the surface (100). At 0% saturation, hue has no effect on the final colour.
- **Brightness**: A relative value within the source being viewed (i.e. a computer monitor or printed document). It is measured in percent, from black (0) to white (100). At 0% brightness, both hue and saturation have no effect on the final generated colour.

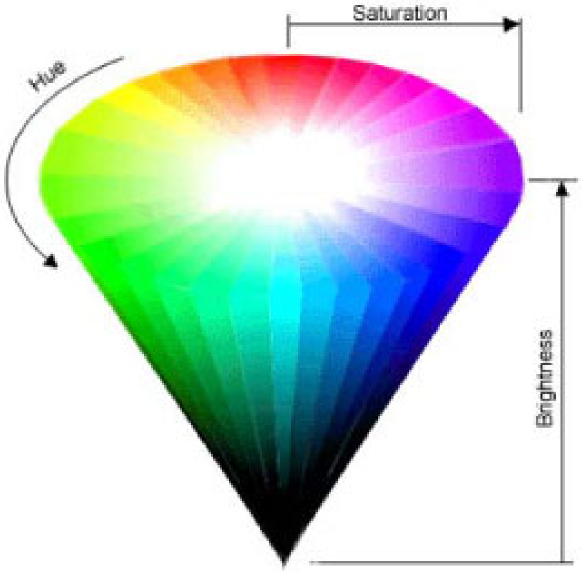

Figure 3.15: *Visual representation of HSB model*[1]

The HSB model suffers from a common limitation of computation colour models - the mapping of brightness over a linear scale, when the eye actually senses brightness approximately logarithmically over a moderate range (Cornsweet, 1970). However, due to the relatively large changes in contrast produced in this

---

[1] from `http://www.tomjewett.com/colors/hsb.html`. Date last accessed: November 6, 2012

experiment, and the simplicity in contrast control the HSB model provides, it is the model used in this study.

3.5.1 Hypothesis

Primary

If the illumination model adds useful depth information to the scene, mean target completion time will decrease in the $Z_w Z_v$ mapping condition. Furthermore, should enough useful depth information be added by the contrast cues, mean target completion time in the $Z_w Z_v$ mapping will not be significantly different to the mean completion time in the $X_w X_v$ and $Y_w Y_v$ mapping conditions.

Secondary

The No-Contrast group is essentially a control group, as used in other experiments. Results are expected to be consistent with other experiments utilising mapping as a variable (see Sections 3.2, 3.3, 3.4). Specifically, mean target completion time in the $Z_w Z_v$ mapping will be significantly greater than in $X_w X_v$ and $Y_w Y_v$ mapping conditions.

3.5.2 Methods

Participants

In total, 28 subjects were recruited from a university cohort. The mean age was 25.74 years ($SD = 6.12$), and the group was 92% male. Participants had self-declared normal or corrected to normal vision, and none were aware of any vision deficiencies. This was confirmed using a visual acuity test, in which all subjects were assessed to have a VA of at least 20/20 at 3 meters. Additionally, the contrast sensitivity of all participants was assessed using a Pelli-Robson contrast sensitivity chart, and confirmed to be within the normal range for their age (1.65-1.95, as described by Elliott et al. (1990); Mäntyjärvi and Laitinen (2001)).

Procedure

The basic task remains similar to that described in Section 3.2, and is described, with differences here. A grey circular disk was shown on the display screen, of the colour HSB$(0, 0, B)$, where the initial B is predetermined ($0 \leq B \leq 100$). In addition to controlling the position of the disk, movement of the subject's dominant foot, in a single dimension, controlled the value of B, with 10mm of movement equating to a change of 1 in B. During the task, thirty targets appeared sequentially, at predetermined locations on screen, coloured with one of 30 predetermined grey (HSB$(0, 0, B)$) levels. The order of grey level presentation was the same for each subject. The subjects were instructed to match the colour of the disk with the colour of the target by moving their foot. In order to familiarise them with the task, subjects initially completed a short training protocol of three targets. During training, the mapping condition was consistent with the assigned experimental condition, and the contrast condition was selected to be the second of the assigned contrast presentation order, to minimise the effect of learning on the experimental outcome. Accurate matching of the disk and target colours was indicated by the screen border changing from red to green, and after the subject had maintained accurate foot position for five seconds, the target and disk B values changed. The main presentation consisted of thirty targets for each contrast condition. Target locations were consistent with previous experiments, and were generated as described in Section 3.2.2. The disk-target contrast level were pre-calculated in a similar way. To minimise fatigue, there was a break of at least three minutes between the training and experimental presentations, and subjects were allowed a rest period of up to five minutes between each contrast condition. To ensure progression through the task, each target was time-limited to sixty seconds, after which the next target was presented. 'Failed' targets were presented again at the end of the sequence.

Design

A mixed design was used, with a between-subjects factor of mapping, and a within-subjects factor of contrast condition. Each subject was randomly assigned to one of four mapping groups: '$X_w X_v$', '$Y_w Y_v$', '$Z_w Z_v$', or '$Z_w Y_v$'. Within each mapping group, subjects completed both the No-Contrast and With-Contrast conditions. The dependant variable is mean time to target completion, defined as the time taken to complete all targets within a mapping*contrast condition divided by the number of presented targets. In order to minimise any learning effects, the presentation order of the contrast condition was randomised for each subject.

Statistics

It is expected that a two-way mixed design ANOVA will show a main effect of mapping, consistent with prior experiments. Contrasts between the four mapping levels are also expected to maintain consistency with prior experiments. Two-way ANOVA analysis will also illuminate a main effect of contrast, should there be a significant difference between the No-contrast and With-contrast conditions. Finally, the ANOVA will demonstrate any interaction between mapping and contrast conditions. Should all hypothesis hold true, it is expected that there will be significant effects in all three of the above, and these will then be further investigated using post-hoc comparisons, with Bonferroni adjusted significance levels as required.

3.5.3 Results

For each subject, the total completion times for both the No-Contrast and With-Contrast conditions were divided by the number of trials in each contrast condition to obtain the mean completion times for each contrast condition. The mean of these times was then calculated across the seven subjects in each of the mapping conditions to obtain a grand mean trial completion time for each mapping-by-contrast condition. The resulting means and standard deviations can be found in Table 3.10.

Condition	Mapping	Mean	Standard Deviation
No Contrast	$X_w X_v$	7.715	0.342
	$Y_w Y_v$	7.653	0.312
	$Z_w Y_v$	8.146	0.431
	$Z_w Z_v$	9.933	0.237
With Contrast	$X_w X_v$	7.601	0.508
	$Y_w Y_v$	7.372	0.237
	$Z_w Y_v$	8.334	0.214
	$Z_w Z_v$	9.883	0.604

Table 3.10: *Means and standard deviations of single-trial completion times (in seconds) for each mapping-by-contrast condition*

The data were tested and found to have met the assumptions of Levene's test of homogeneity of variance for the between-subjects independent variable (mapping, Table 3.11). Both levels of the within-subjects factor are shown in Table 3.11 as the homogeneity of variance across the levels of the mapping factor were tested within both levels of the contrast factor. The non-significant p-values indicate that the distributions of variance across all levels of mapping within both levels of contrast do not differ from a normal distribution.

Condition	F	df1	df2	Significance
No Contrast	0.242	3	24	.866
Contrast	2.315	3	24	.101

Table 3.11: *Levene's test of equality of variances for the between-subjects mapping factor*

ANOVA Results

Results of a two-way mixed ANOVA showed that the main effect of contrast was non-significant, $F(1, 24) = 0.34$, $p = .567$, (Figure 3.16). Consistent with previous experiments, there was a main effect of mapping, $F(3, 24) = 134.81$, $p < .001$, as shown in Figure 3.17. There was no mapping-by-contrast interaction effect, $F(3, 24) = 0.78$, $p = .518$, (Figure 3.18).

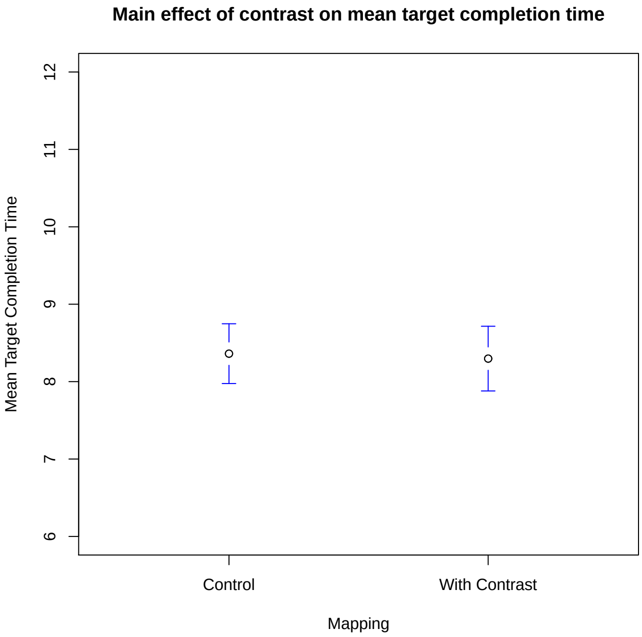

Figure 3.16: *Main effect of contrast on mean target completion times (in seconds). Error bars represent 95% confidence intervals.*

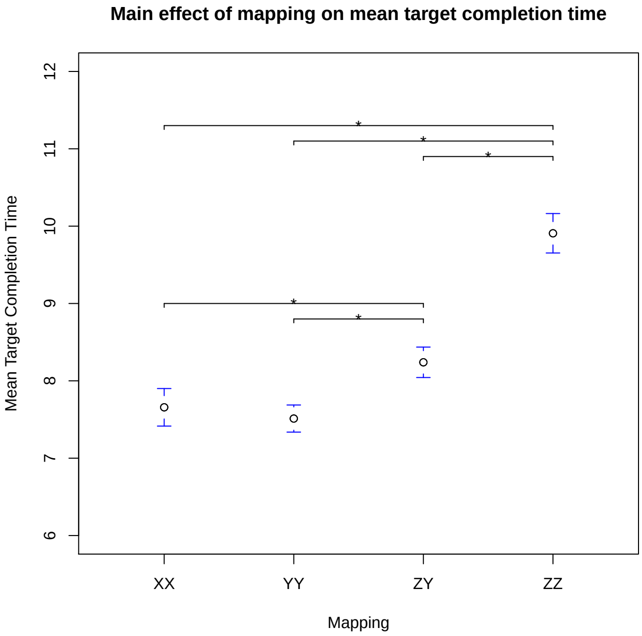

Figure 3.17: *Main effect of mapping on mean target completion times (in seconds). Error bars represent 95% confidence intervals. ‘ * ’ indicates significance at* $p < .05$.

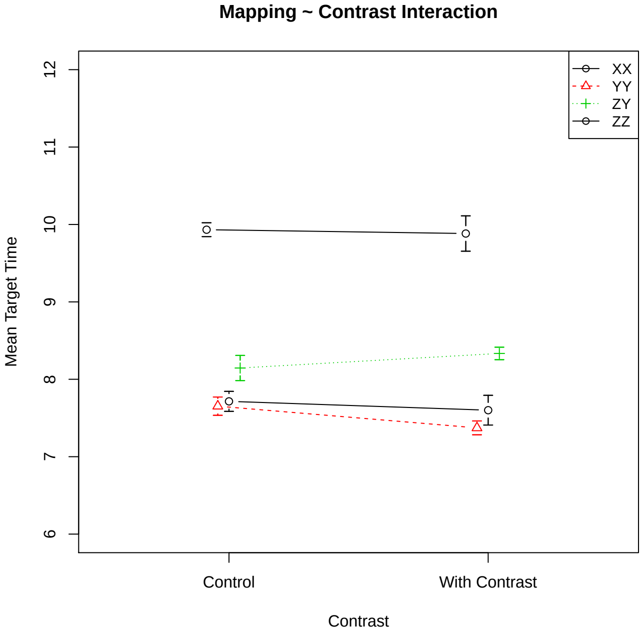

Figure 3.18: *Mapping-by-contrast interaction effect on mean target completion times (in seconds). Error bars represent 95% confidence intervals.*

Mapping Pairwise Comparisons

As the above analysis indicates a significant main effect of mapping, this was investigated further using pairwise comparisons across each contrast condition (Table 3.12. Comparisons between mappings are consistent with previous experiments. Mean completion time in Z_wZ_v is greater than all other mappings ($p < .001$, compared to X_wX_v, Y_wY_v, Z_wY_v). Mean completion time in Z_wY_v was also greater than X_wX_v ($p = .025$) and Y_wY_v ($p = .012$).

Contrast Condition	Mapping (I)	Mapping (J)	Mean Difference (I-J)	Significance
No-contrast	X_wX_v	Y_wY_v	0.062	.734
		Z_wY_v	-0.431	.025
		Z_wZ_v	-2.217	.000
	Y_wY_v	X_wX_v	-0.062	.734
		Z_wY_v	-0.493	.012
		Z_wZ_v	-2.280	.000
	Z_wY_v	X_wX_v	0.431	.025
		Y_wY_v	0.493	.012
		Z_wZ_v	-1.787	.000
	Z_wZ_v	X_wX_v	2.217	.000
		Y_wY_v	2.280	.000
		Z_wY_v	1.787	.000
With-contrast	X_wX_v	Y_wY_v	0.229	.325
		Z_wY_v	-0.733	0.004
		Z_wZ_v	-2.282	.000
	Y_wY_v	X_wX_v	-0.229	.325
		Z_wY_v	-0.962	.000
		Z_wZ_v	-2.511	.000
	Z_wY_v	X_wX_v	0.733	.004
		Y_wY_v	0.962	.000
		Z_wZ_v	-1.549	.000
	Z_wZ_v	X_wX_v	2.282	.000
		Y_wY_v	2.511	.000
		Z_wY_v	1.549	.000

Table 3.12: *Mapping-by-contrast pairwise comparisons showing differences in mean trial completion times (in seconds). Significance values shown are Bonferroni-corrected and can therefore be interpreted as significant at $p < 0.05$.*

Results Summary

The results of the current experiment show that isolating a contrast cue and adding that to a minimal scene did not reduce the mean target completion time. Consistent with previous experiments, there was a significant effect of dimensional World-to-Screen mapping on task performance, and this effect did not differ between the two contrast levels.

3.6 Scaling

In the real world, the effect of perspective leads to objects appearing smaller with increasing distance from the observer, due to the angle subtended by the object reducing. This effect can be simulated within virtual environments through the utilisation of a scaling effect along a consistent, defined one-point perspective viewpoint.

3.6.1 Hypothesis

Primary

If scaling adds useful depth information to the scene, mean target completion time will decrease in the $Z_w Z_v$ mapping condition. Furthermore, should enough useful depth information be added by the scaling, mean target completion time in the $Z_w Z_v$ mapping will not be significantly different to the mean completion time in the $X_w X_v$ and $Y_w Y_v$ mapping conditions.

Secondary

Scaling values are varied parametrically above and below 1. The condition in which the scaling value is equal to 1 is essentially a control condition, where subjects complete the same task as that used in the previous experiments. Results are expected to be consistent with other experiments utilising mapping as a variable (see Sections 3.2, 3.3, 3.4, 3.5). Specifically, mean target completion time in the $Z_w Z_v$ mapping will be significantly greater than in $X_w X_v$ and $Y_w Y_v$ mapping conditions.

3.6.2 Methods

Participants

In total, 28 subjects were recruited from a university cohort. The mean age was 27.87 years ($SD = 5.43$), and the group was 86% male. Participants had self-declared normal or corrected to normal vision, and none reported any vision deficiencies. This was confirmed using a visual acuity test, in which all subjects were assessed to have a VA of 20/20 at 3 meters.

Design

A mixed design was used, with a between-subjects factor of mapping, and a within-subjects factor of scaling condition. Each subject was randomly assigned to one of four mapping groups: 'X_wX_v', 'Y_wY_v', 'Z_wZ_v', or 'Z_wY_v'. Within each mapping group, subjects completed the full range of scaling conditions, in which their movement was scaled by a gain factor ranging from 0.25 to 2.0. This is described in more detail below. The dependant variable is mean time to target completion, defined as the time taken to complete all targets within a mapping*scaling condition divided by the number of presented targets. In order to minimise any learning effects, the presentation order of the scaling condition was randomised for each subject.

Procedure

The basic task remains similar to that described in Section 3.2, and is described, with differences here. Eight scaling conditions were tested, as shown in Table 3.13.

Scaling Factor:	0.25	0.5	0.75	1	1.25	1.5	1.75	2

Table 3.13: *Scaling factors*

A short training protocol was administered, in which each subject completed three targets within their assigned mapping condition. The scaling condition for each target was randomly assigned, to acclimatise subjects to differing scaling conditions, while minimising learning effects. The main presentation consisted of thirty targets at locations consistent with those described in Section 3.2.2. To minimise fatigue, there was a break of at least three minutes between the training and experimental presentations, and subjects were allowed a rest period of up to five minutes between each scaling condition. To ensure progression through the task, each target was time-limited to sixty seconds, after which the next target was presented. 'Failed' targets were presented again at the end of the sequence.

Statistics

It is expected that a two-way mixed design ANOVA will show a main effect of mapping, consistent with prior experiments. Contrasts between the four mapping levels are also expected to maintain consistency with prior experiments. Two-way ANOVA analysis will also illuminate a main effect of scaling, should there be a significant difference between the No-scaling and With-scaling conditions. Finally, the ANOVA will demonstrate any interaction between mapping and scaling conditions. Should all hypothesis hold true, it is expected that there will be significant effects in all three of the above, and these will then be further investigated using post-hoc comparisons, with Bonferroni adjusted significance levels as required.

3.6.3 Scaling

For each subject, the total completion times within each scaling condition were divided by the number of targets in each scaling condition to obtain the mean single-trial completion times for each scaling condition. The mean of these times was then calculated across all seven subjects in each of the mapping conditions to obtain a grand mean single-trial completion time for each mapping-by-scaling condition. The resulting means and standard deviations are shown in Table 3.14.

The data were tested and found to have met the assumptions of Mauchly's test of sphericity ($p = .388$) for the within-subjects scaling factor. The data also met the assumptions of Levene's test of equality of variances for the between-subjects factor (mapping, Table 3.15). All levels of the within-subjects factor are shown in Table 3.15 as the homogeneity of variance across the levels of the mapping factor were tested within each level of scaling. The non-significant p-values indicate that the distributions of variance across all levels of mapping within each level of scaling do not differ from a normal distribution.

ANOVA Results

A two-way mixed ANOVA showed that there was a within-subjects main effect of scaling, $F(7, 168) = 746.78$, $p < .001$, and a between-subjects main effect of mapping, $F(3, 24) = 297.69$, $p < .001$. There was also a mapping-by-scaling interaction effect, $F(3, 24) = 120.44$, $p < .001$.

Scaling Factor	Mapping	Mean	Standard Deviation
0.25	$X_w X_v$	13.37	0.54
	$Y_w Y_v$	12.99	0.57
	$Z_w Y_v$	13.80	0.38
	$Z_w Z_v$	14.68	0.39
0.5	$X_w X_v$	10.85	0.81
	$Y_w Y_v$	10.49	0.25
	$Z_w Y_v$	11.45	0.31
	$Z_w Z_v$	11.50	0.45
0.75	$X_w X_v$	9.01	0.25
	$Y_w Y_v$	9.08	0.36
	$Z_w Y_v$	9.79	0.22
	$Z_w Z_v$	9.81	0.21
1	$X_w X_v$	7.49	0.42
	$Y_w Y_v$	7.48	0.38
	$Z_w Y_v$	8.05	0.31
	$Z_w Z_v$	10.13	0.27
1.25	$X_w X_v$	9.27	0.44
	$Y_w Y_v$	8.81	0.33
	$Z_w Y_v$	9.82	0.46
	$Z_w Z_v$	9.08	0.44
1.5	$X_w X_v$	10.95	0.50
	$Y_w Y_v$	10.61	0.44
	$Z_w Y_v$	11.75	0.29
	$Z_w Z_v$	6.42	0.32
1.75	$X_w X_v$	7.20	0.60
	$Y_w Y_v$	5.57	0.61
	$Z_w Y_v$	5.93	0.41
	$Z_w Z_v$	7.79	0.41
2	$X_w X_v$	15.70	0.39
	$Y_w Y_v$	15.62	0.47
	$Z_w Y_v$	17.00	0.31
	$Z_w Z_v$	9.34	0.30

Table 3.14: *Means and standard deviations of single-trial completion times (in seconds) for each mapping-by-scaling condition*

Scaling factor Effect

As Table 3.16 summarises, planned comparisons showed that each level of the scaling factor was significantly different to the control condition (scaling = 1).

Compared to a scaling factor of 1 (no adjustment in scaling), when mapping is not considered, any adjustment in the scaling factor increases the mean target completion time. This main effect of scaling can be seen in Figure 3.19.

Mapping

As seen in Table 3.17, planned comparisons show that the effect of mapping is consistent with previous experiments for $X_w X_v$, $Y_w Y_v$ and $Z_w Y_v$ mappings. Figure 3.20 shows that scaling significantly reduces the mean target completion time in the $Z_w Z_v$ mapping, and this is further investigated in Section 3.6.3.

Mapping-by-Scaling Interaction

As can be seen from Figures 3.21a, 3.21b, 3.21c, the mean target completion time generally increases when movement is scaled by any factor other than 1. However, Figure 3.21d shows that completion time in the $Z_w Z_v$ mapping decreases when the scaling factor is increased above 1. In the $Z_w Z_v$ mapping, mean completion time appears to be lowest around a scaling factor of 1.5, after which it increases, though the mean trial completion time is still significantly higher at a scaling factor of 1 than it is at the highest measured scaling factor of 2.

Figure 3.22 also markedly demonstrates the decrease in mean target completion time in the $Z_w Z_v$ mapping for scaling factors above 1.

Scaling Factor	F	df1	df2	Significance
0.25	.292	3	24	.831
0.50	1.650	3	24	.204
0.75	.388	3	24	.763
1.00	.272	3	24	.845
1.25	.397	3	24	.756
1.50	.905	3	24	.453
1.75	.339	3	24	.797
2.00	.201	3	24	.894

Table 3.15: *Levene's test of equality of variances for the between-subjects mapping factor*

Scaling Factor	Mean Difference	Significance
0.25	5.45	$< .001$
0.50	2.81	$< .001$
0.75	1.16	$< .001$
1.25	0.95	$< .001$
1.50	1.67	$< .001$
1.75	3.52	$< .001$
2.00	6.14	$< .001$

Table 3.16: *Significant differences in mean target completion times (in seconds) between each level of scaling and the control condition (scaling = 1)*

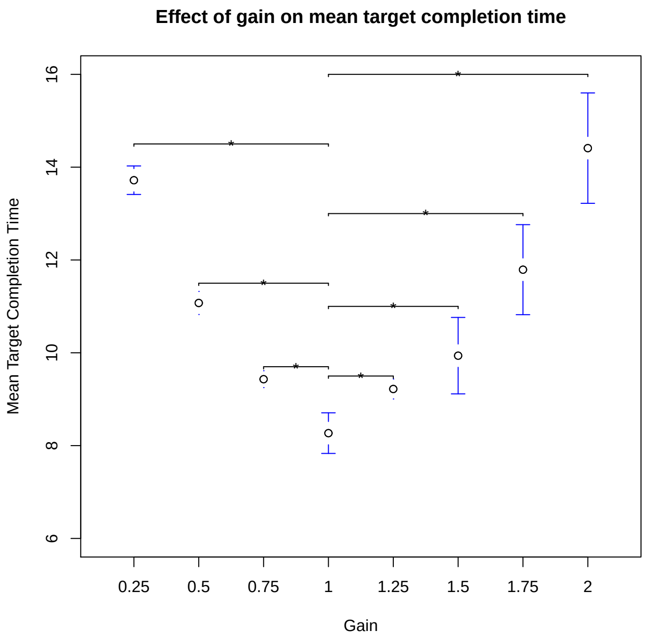

Figure 3.19: *Effect of scaling factor on mean target completion time (in seconds) irrespective of mapping. Error bars represent 95% confidence intervals. ' * ' indicates significance at* $p < .05$.

Mapping (I)	Mapping (J)	Mean Difference (I-J)	Significance
X_wX_v	Y_wY_v	0.16	.190
	Z_wY_v	-0.78	$< .001$
	Z_wZ_v	1.34	$< .001$
Y_wY_v	Z_wY_v	-0.94	$< .001$
	Z_wZ_v	1.18	$< .001$
Z_wY_v	Z_wZ_v	2.12	$< .001$

Table 3.17: *Significant differences in mean target completion times (in seconds) between all levels of mapping irrespective of scaling.*

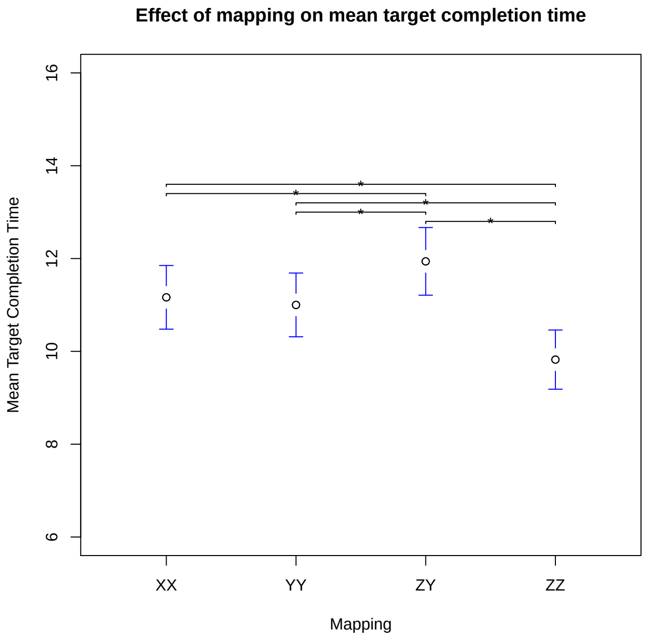

Figure 3.20: *Effect of mapping on mean target completion times (in seconds) irrespective of scaling factor. Error bars represent 95% confidence intervals. ‘ * ’ indicates significance at* $p < .05$.

Results Summary

Below a scaling factor of 1 (where "1" indicates no scaling), the mean target completion time significantly increased for all mappings. That is, performance suffered in all mappings with scaling values below 1. When the scaling factor was greater than 1, the mean completion times increased for mappings X_wX_v, Y_wY_v, and Z_wY_v. In the X_wX_v, Y_wY_v and Z_wY_v mappings, performance was best with no change in scaling. However, in the Z_wZ_v mapping, the mean trial completion times were significantly lower when a scaling factor of above 1 was applied, with the lowest mean trial completion time (i.e. best performance) occurring in the 1.5 scaling condition.

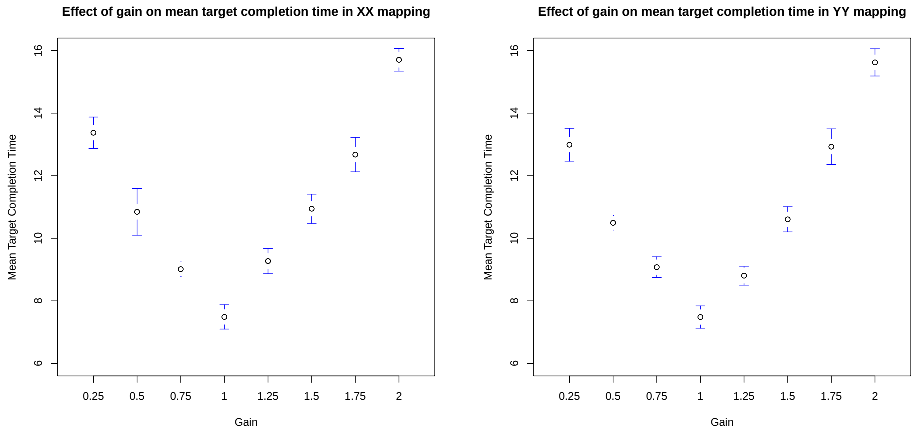

(a) X_wX_v **(b)** Y_wY_v

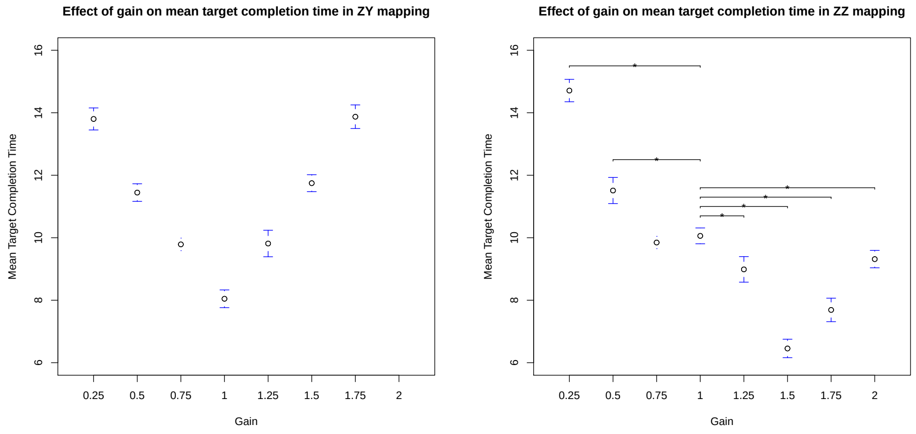

(c) Z_wY_v **(d)** Z_wZ_v

Figure 3.21: *Effect of scaling factor (gain value) on mean trial completion times (in seconds) within each mapping. Error bars represent 95% confidence intervals. Tiny error bars not shown. To maintain statistical power, significance of differences between levels of scaling was only calculated for the* Z_wZ_v *mapping. ‘ * ’ indicates significance at* $p < .05$.

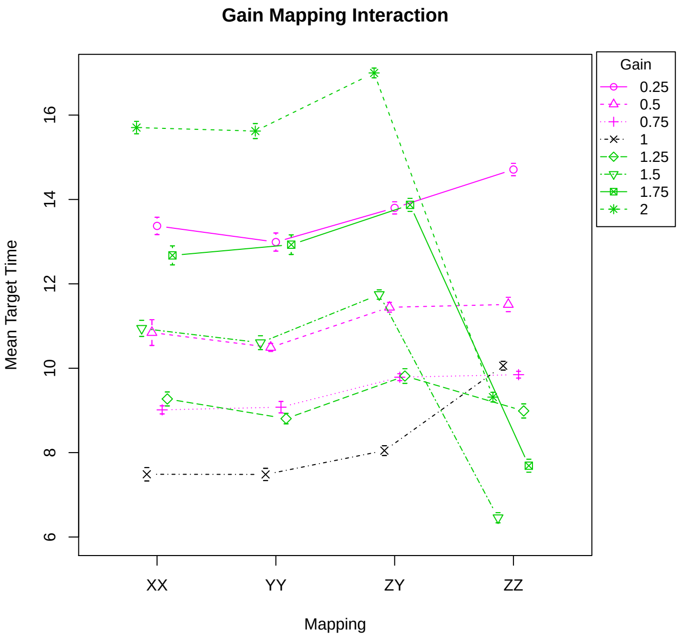

Figure 3.22: *Mapping-by-scaling interaction effect on mean trial completion times (in seconds). Scaling factors above and below 1 shown in green and purple respectively.*

3.7 Obstacle Avoidance

In order to test the model derived from the results of the scaling study described in Section 3.6 and discussed in Section 4.5, a realistic, unconstrained task of the type utilised in gait investigation and rehabilitation can be used. Obstacle avoidance tasks have been repeatedly demonstrated to provide an adequate paradigm to study the human stepping pattern under controlled experimental conditions (McIntosh et al., 2004; Katz et al., 2005; Holden, 2005; Den Otter et al., 2005; Gérin-Lajoie et al., 2008), and as such, a customised obstacle avoidance task was designed and implemented for this study.

3.7.1 Hypothesis

Based on the predications of the model under test, task performance will optimal when a scaling factor of 1.5 is applied to real-world movement in the sagittal (Z) plane.

3.7.2 Methods

Participants

In total, 18 subjects were recruited from a university cohort. The mean age was 27.31 years ($SD = 5.74$), and the group was 83% male. Participants had self-declared normal or corrected to normal vision, and none reported any vision deficiencies. This was confirmed using a visual acuity test, in which all subjects were assessed to have a VA of at least 20/20 at 3 meters.

Design

A within-subjects design was used, with every subject completing all levels of the scaling factor; 1 (no scaling), 1.5, and 1.75. Unlike the previous experiments, the dependent variable in this task is the average number of boxes correctly contacted ('kicked') and avoided over three trials in each scaling condition.

Procedure

The task consisted of a virtual corridor populated with a number of three-dimensional target boxes, similar to those in Figure 3.23. In each trial, thirty boxes were distributed evenly along the corridor, fifteen of which were coloured purple, and fifteen were coloured yellow. Subjects were instructed to ‘kick’, or otherwise contact the yellow boxes with their avatar, controlled using the motion capture system described in previous experiments. They were also instructed to avoid the purple boxes, and were told that their final score was the sum of correctly ‘kicked’ and avoided boxes. Subjects started the task on a stationary treadmill, positioned one meter from the display screen. On task commencement, the treadmill speed was smoothly increased to a comfortable walking pace of 3 mph.

Subjects completed three consecutive trials in each level of the scaling factor. The order of scaling level presentation was randomised to minimise order effects. Subjects were allowed to take breaks of up to 2 minutes between the trial sets.

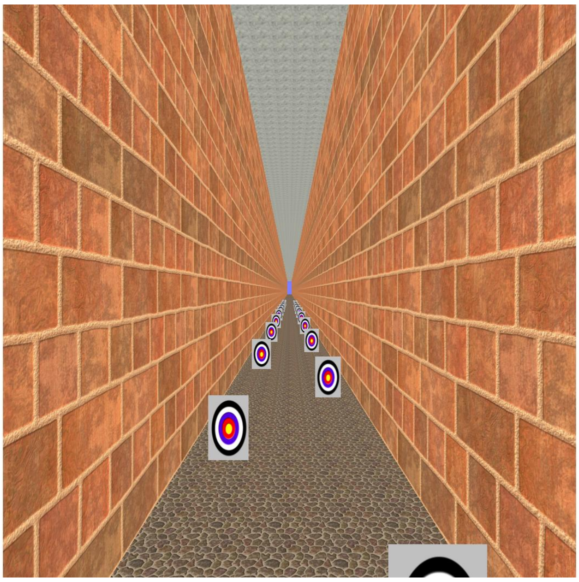

Figure 3.23: *The obstacle avoidance task used to test the proposed scaling model.*

Statistics

A one-way repeated-measures ANOVA will be used to determine the effect of Z-dimensional scaling on the task score (total number of boxes successfully contacted and avoided). If the effect of scaling on task score is found to be significant, the ANOVA result will be followed up with multiple comparisons to determine which sets of scaling conditions differed significantly.

Given the results of the scaling study (see Section 3.6.3), it is expected that the Z_v dimension scaling factor will have a significant effect on performance. If the hypothesis is correct, then the task score will be significantly higher in the 1.5 scaling condition relative to both the 1 and 1.75 scaling conditions.

3.7.3 Results

For each subject, the total number of correct ‘kicks’ (successful avatar contact with yellow boxes) and correct misses (successful avatar avoidance of purple boxes) for each trial was divided by the total number of trials completed within that Z_v-scaling condition. This resulted in the mean task scores within each scaling condition for each subject. The resulting task scores were then averaged across all subjects, producing a grand mean task score for each grating condition. The means and standard deviations of the task scores for each level of the scaling factor are summarised in Table 3.18.

Z_v Scaling Condition	Mean	Standard Deviation
No Scaling	18.46	1.35
1.5 Scaling	27.26	1.12
1.75 Scaling	24.32	1.36

Table 3.18: *Means and standard deviations of task scores (mean number of boxes correctly contacted/avoided per trial) for each scaling condition. The maximum possible score in the task was 30.*

The value of Mauchly’s test for sphericity was non-significant ($p = .497$). This result indicates that the data does not violated the assumption of sphericity and has met the criteria for a one-way repeated-measures ANOVA.

ANOVA Results

A one-way repeated-measures ANOVA revealed that there was a significant effect of scaling on the task score, $F(2, 34) = 190.69$, $p < .001$.

Multiple Comparisons Between Scalings

Multiple comparisons showed that there are significant differences between all scaling conditions, and these are summarised in Table 3.19. The mean difference between the 1.5 Scaling and the No-Scaling conditions is 8.80 boxes. This means that on average, subjects successfully navigated (i.e. correctly kicked or avoided) about 8-9 more boxes out of 30 boxes total when the Z_v dimension was scaled by a factor of 1.5 compared to a factor of 1. The difference in the mean task scores between the 1.75 Scaling and No-Scaling conditions was of the same sign but a lesser magnitude. Finally, the comparison between the mean tasks scores for the 1.5 and 1.75 Scaling conditions shows an advantage of about 3 boxes in the 1.5 Scaling condition.

Z_v Scaling (I)	Z_v Scaling (J)	Mean Difference (I-J)	Significance
None	1.5	-8.80	$< .001$
None	1.75	-5.85	$< .001$
1.5	1.75	2.94	$< .001$

Table 3.19: *Comparisons of mean task scores between each pair of scaling levels. All differences between scaling conditions are significant at the Bonferroni-corrected alpha value of* .017.

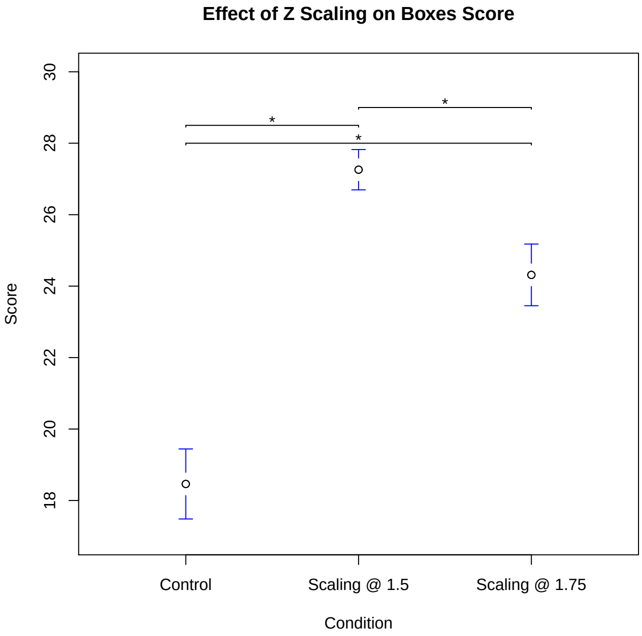

Figure 3.24: *Effect of scaling on the obstacle avoidance task score. ‘ * ’ indicates significance at* $p < .05$. *Error bars represent 95% confidence intervals.*

Figure 3.24 shows the significant increase in task performance when a scaling factor is applied to the Z_v dimension. The highest mean task score occurred when the scaling factor at 1.5.

Results Summary

Results of the one-way repeated-measures ANOVA showed that performance in the task differed depending on the scaling factor. Post-hoc mulitple comparison tests revealed that task scores in all three levels of the scaling factor were significantly different from one another. The mean task score was highest when the Z_v dimension was scaled to 1.5, followed by the 1.75 scaling. Mean task scores were lowest in the No-Scaling (scaling = 1) condition.

Chapter 4

Discussion

This chapter will discuss the results of each experiment described in Chapter 3, before pulling them all together and examining the primary statement of this thesis.

4.1 Mapping

The superior performance of “straight” mappings X_wX_v and Y_wY_v, compared to Z_wZ_v demonstrates that the initial hypothesis of reduced performance in Z_wZ_v was correct. It is theorised that this is due to the increased visuospatial processing required when Z_v scaling is used to indicate subject movement in that plane and is in keeping with Fitts’ law. The similarity of X_wX_v and Y_wY_v performances supports this - the onscreen target remains the same size in each of these mappings, thereby requiring the same level of visuospatial processing. Reducing the degradation of performance in the Z_v plane is of particular importance when utilising virtual reality tasks for the assessment of motor deficits, and also in virtual rehabilitation tasks, where visual feedback is provided to a possibly cognitively impaired patient and is relied on as a primary source of motion information.

Looking at the World Z_w level, we see that Z_wY_v is the most optimal mapping, followed by Z_wX_v, then Z_wZ_v. While Z_wY_v performance still lags behind X_wX_v and Y_wY_v, it is an improvement over the alternate mappings. It is proposed that this is due reduced visuospatial processing required in decoding this mapping, combined with the motion cue similarities of a real-world subject looking down at their own feet (feet move up-down relative to the field of view), and a subject looking forwards at the screen, with the forward-backward movement of their feet translated into up-down motion onscreen.

4.2 Grating

The discovery of the significant main effect of grating on mean target completion time shows that working within a grating paradigm does improve spatial accuracy, as hypothesised. The specific, significant improvement in distinct grating-mapping conditions (Vertical improves X_wX_v, Horizontal improves Y_wY_v and Z_wY_v), and lack thereof in other grating-mapping conditions, leads to the conclusion that the addition of *relevant* visuospatial information improves motion accuracy. The lack of any improvement in the Z_wZ_v mapping suggests that the addition of perspective alone may not improve spatial accuracy, and also demonstrates that solving the dimensional mapping problem described in Section 3.2 is non trivial.

This experimental work is also consistent with the previous experiment, again showing that the most optimal mappings are X_wX_v, Y_wY_v and Z_wY_v. Further work is required to investigate methods of improving performance in the Z_wZ_v mapping to a level comparable with the X_wX_v and Y_wY_v mappings.

These, and previous results, show that utilising a fully three-dimensional environment for the enhancement of spatial motor control may not provide the most optimal visualisation. In a stroke rehabilitation setting, where elderly patients are unlikely to have any experience with motion controlled interfaces and may also be suffering cognitive and/or visual impairments in addition to motor control deterioration, any possible additional cognitive load created by a sub-optimal motion controlled system is likely to negate the desired positive effects of the rehabilitation equipment. Interference between gait and cognitive tasks is well demonstrated by Haggard et al. (2000), and it is therefore highly desirable to continue to investigate approaches to improving three-dimensional spatial accuracy within a three-dimensional virtual rehabilitation environment, while also further investigating the efficacy of two-dimensional virtual tasks for the same rehabilitation purposes.

4.3 Stereoscopic 3D

This study demonstrated that anaglyph 3D does not have any effect on improving speed and accuracy performance in the simple virtual task. Contrary to the hypothesis, the mean completion time in the with-3D condition was not found to be significantly different to the mean completion time in the control condition. This is consistent with similar studies, including Hanna et al. (1998), in which two-dimensional and three-dimensional imaging in elective laparoscopic surgery was compared, in which no difference in median completion time was found. Liu et al. (1993) also found that they could eliminate stereopsis from their display system with no negative effect, incidentally, when attempting to decrease the amount of processing their system had to undertake, and Peli (1998) finds no difference in functional changes in the visual system after the use of binocular head-mounted displays and mono CRT monitors. Perhaps indicative as to the difficulty in quantifying performance in this field of study however, there is some evidence that the addition of stereoscopic cues can aid in the performance of virtual tasks requiring depth perception. Rosenberg (1993) reports a tenfold increase in performance in their virtual task when using stereoscopic cues, and he further goes on to find that stereoscopic displays increase accuracy in elevation estimation, though not in estimating relative azimuth direction (Barfield and Rosenberg, 1995).

While it can not be said that there is widespread consensus on the efficacy of stereoscopic displays in improving performance in tasks requiring depth perception, two primary arguments can be made for not discounting stereoscopic 3D as a source of additional depth information. It may be the case that there was too little information in the virtual scene used in this task, and a scene containing more objects might provide more opportunities for binocular parallax. Additionally, the effect of stereopsis may enhance the effect of other depth cues, and not produce a large effect when tested alone, as in this task. It is also possible that it may take time to adjust to the effect of anaglyph 3D, but that increased performance, as compared to a no-3D condition, is obtainable after adjustment.

Though evidence into the benefits of stereoscopic displays is mixed, Stone et al. (2012) find evidence of harmful effects. They use the Randot depth perception test, as well as a number of practical tests to assess subjects depth perception before and after watching 4 hours and 45 minutes of 3D television, finding significant degradation after 3D television exposure. Their study does suffer from a number of shortcomings however, primarily a lack of comparison with exposure to a standard television for a similar period of time. The potential detrimental effects of screen viewing on vision, for example, are well known

(Kumar, 2006; Pavithra, 2011), and as such, it is hard to draw conclusions as firm as the claims of the authors. However, in the largest study of visual stress from 3D viewing, Atallah et al. (2012) find that greater than 40% of observers suffer a range of effects from 3D stereoscopic displays under typical viewing conditions.

Within this study it is of additional interest to note that the effect of mapping is consistent with previous experiments. Mean target completion time is greatest in the Z_wZ_v mapping, and is significantly greater than the X_wX_v and Y_wY_v mappings.

In summary, the results of this study, and the lack of strong evidence in the field, leads to the conclusion that the use of stereo in the Stromohab system can be discounted for at least the initial use cases.

4.4 Contrast

This study demonstrated that a using a simple illumination model did not have an effect on accuracy and performance as measured by the simple virtual task. Again, contrary to the hypothesis, the mean completion time in the with-contrast condition was not found to be significantly different to the mean completion time in the no-contrast condition (Section 3.5.3). This result, while unexpected, does indicate that this simple lighting model does not provide useful depth information in the context of the Stromohab system. While lighting models are not the subject of this thesis, it is worth nothing that there are many highly complex lighting and shading models available to virtual environment designers, though the objective evaluation of these in contexts similar to this study are difficult to find. Shioiri et al. (2012) do study the effects of colour and luminence on two-dimensional and three-dimensional motion, and propose a new low-level neurological model responsible for motion integration within the brain, while Bex et al. (2007) have questioned the relevance of studying contrast using the prevaling sinsusoidal grating model to the way in which we perceive the natural world. Theses studies, and others in their vein, are undoubtably useful in adding to our understanding of perception and the the human brain. Additionally, when combined with high-end computer graphics technology, they can be of some use in designing and improving the realism of virtual environments. Complex lighting models however, are invariably used to light complex virtual scenes, and many of the effects produced by high-end graphical hardware would likely not be seen in a simplistic, stripped-down environment as used in this virtual task, effectively nullifying the features of more powerful illumination models. This study therefore, discards the simple illumination model for the purposes of this task, though as shown by Slater et al. (1995) and Haggard et al. (2000), more complex illumination techniques are likely to provide useful depth information when used in more complex virtual environments.

4.5 Scaling

This study shows that applying a scaling factor to movement in the transverse (Z_v) plane affects depth plane accuracy as tested by the simple virtual task used through this thesis. Consistent with the hypothesis, there is an improvement when Z_w plane movement is 'magnified', and furthermore, optimal scaling factors are found, which produce a performance improvement large enough to eliminate the previously significant deficit in the transverse plane, as compared to coronal (X_w) and sagittal (Y_w) planes of movement. This improvement is most marked at a scaling factor of 1.5, after which it decreases, remaining however significantly better than a scaling factor of 1 at the highest measured factor of 2. Applying any scaling to movement in the X_w and Y_w planes leads to a highly similar, symmetrical degradation in performance, for scaling factors of both above and below one.

This result forms a significant part of the contribution to knowledge of this thesis. The problem of distance compression, or underestimation in virtual environments has been observed in many studies (Sahm et al., 2005; Frenz et al., 2007; Grechkin et al., 2010). Some attempts to resolve the cause of this distance underestimation have been made, with most studies eliminating their proposed hypotheses. Witmer and Sadowski (1998) find that the magnitude of relative errors in a virtual environment is twice that in the real world, and suggest that the performance deficit is due to a poor binocular disparity cues, or a distortion of pictorial depth cues. However, as discussed in Section 4.3, binocular cues have not yet been conclusively shown to improve depth perception in virtual environments, and their second proposal is somewhat countered by Thompson and Willemsen (2004), who investigate whether or not the quality of computer graphics is the cause of the issue, using a walking task within environments of varying graphical realism. They conclude that distance underestimation is not caused by a lack of realistic graphical rendering. Knapp and Loomis (2004) look at the issue of the limited field of view provided by head-mounted displays, and conclude that this is not the cause of distance underestimation in virtual environments.

It is worth noting at this point, that as is reflected in the above studies, the notion of distance compression in a virtual environment has primarily come from observations of exocentric movements, defined here as the tendency of subjects to underestimate the absolute (point A to point B) distance they have moved through an environment, not the distance they are moving a limb, as in this study. The distance a subject moves a limb, relative to themselves, is defined as the *egocentric* distance, and is discussed below.

Additionally, as far as is known, the virtual task control paradigm used throughout this thesis is unique, with the closest studies being those utilising upper-body reaching type tasks to explore depth perception. Furthermore, reaching as a systematic measure of depth perception has only been used in the preceding decade or so, other than a study by Foley and Held (1972), which followed earlier prism-adaptation work involving the same author (Held and Gottlieb, 1958; White et al., 1964; Held, 1968), and a number of studies on depth direction perception (Welch, 1978; Bingham and Romack, 1999). These studies find that, without relative distance information (i.e. when the subject could not see a representation of themselves within the virtual environment), there is a tendency to overreach when aiming for a target, though also find that, over time, adaptation and correction for this error is possible when feedback is supplied.

One extensive study with which it may be possible to compare the current paradigm is that of Bingham et al. (2001), in which the authors manipulate accommodation, occlusion, and disparity matching cues in real and virtual environments in order to study their effects on space perception. Their results compare space perception in real and virtual environments, and find that subjects overestimated egocentric distances in the virtual environment, relative to the real world, when no relative distance information was available. They do also investigate reaching when relative distance information is available, again comparing results to the real world. Though their data is comparatively analysed, they do note that their results point to a tendency to overestimate depth, egocentric distance and 3D size in a monocular condition, and underestimate egocentric distance and 3D size in a binocular condition,

Overall, two main conclusions can be drawn from these studies:

1. The absolute distance from the observer to an object (*exocentric*, as defined at the beginning of this section), is generally underestimated in virtual environments.
2. The relative distance the observer needs to move through, in order to reach an object (*egocentric*, as defined previously), is generally overestimated in virtual environments.

Using the above findings, we can thus start developing a model to explain the results of this study. Figure 4.1 shows our initial model, containing a virtual object, represented by the cube, and the avatar, screen and real-world position of the subject's viewpoint.

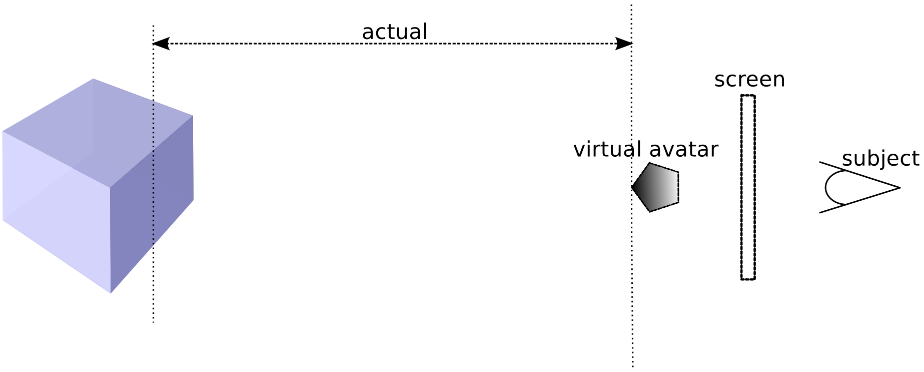

Figure 4.1: *Representation of object and virtual environment.*

Figure 4.2 incorporates item 1 from above into the model, showing the absolute distance underestimation present in virtual environments.

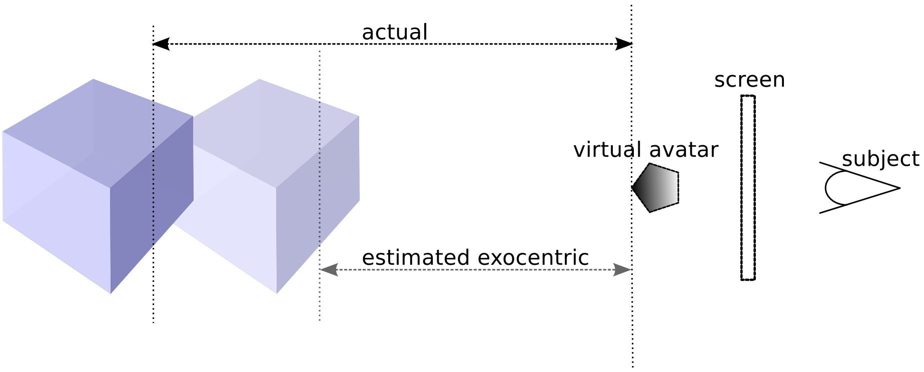

Figure 4.2: *Representation of object and virtual environment, including underestimated exocentric distance.*

A simple model is now apparent:

$$a = x * s \tag{4.1}$$

Where a = the actual position of the virtual object, x = position error caused by exocentric underestimation, and s is the experimental scaling factor. Note that x may be non-linear; this is discussed further below.

This model however, predicts a symmetric response around an optimal value of s. While this approximation may be valid for X_wX_v and Y_wY_v mappings, the experimental results for $Z_w Z_v$ clearly show that this is not the case, and that there are at least two different mechanisms affecting task performance (Figure 3.21d) in the $Z_w Z_v$ mapping.

Figure 4.3 illustrates item 2, egocentric distance overestimation.

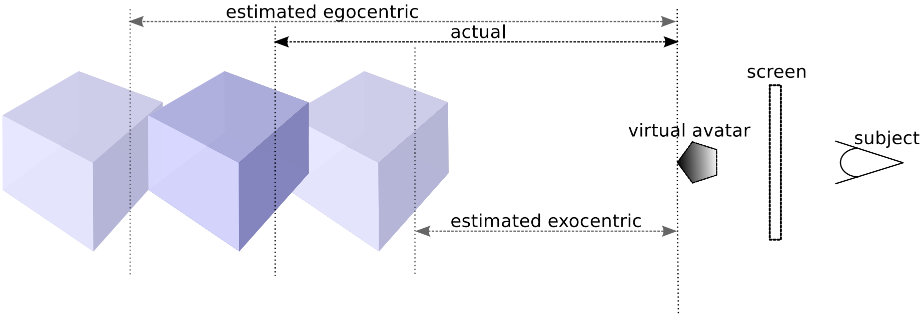

Figure 4.3: *Representation of object and virtual environment, including overestimated exocentric and underestimated exocentric distance.*

The study results (Section 3.5.3) show peak performance when the experimental scaling factor is equal to 1.5, and more generally, performance is consistently improved when the experimental scaling factor is greater than 1, and degraded when the scaling factor is less than 1. We thus deduce that exocentric underestimation has a greater effect than egocentric overestimation, and therefore:

$$\begin{gathered} a = (g - x) * s \\ \text{subject to: } x \geq g,\ s > 0 \end{gathered} \tag{4.2}$$

Where g = the estimated egocentric distance.

Breaking this assertion down, we see that the process for accurate movement to a position in virtual space, using a physically mapped interface can be summarised as follows:

Before movement commences, due to exocentric underestimation, the absolute distance to the virtual target is perceived to be smaller than it is. In a reaching task, the distance the subject is required to move their limb in order to reach the perceived target distance is overestimated, somewhat but not completely compensating for the initial underestimation. The combination of these two factors leads to an overall error, which can be compensated for by applying a positive scaling factor to a subject's sagittal movement.

The results of this study show an optimal scaling factor of 1.5, and our next study tests and validates this assertion.

4.6 Obstacle Avoidance

The observations and conclusions of Section 4.5, are confirmed, demonstrating the proposed practical model is applicable within a three dimensional virtual task of the type likely to be designed and used in applied patient rehabilitation. Consistent with the the hypothesis and the previous study, as compared to the control group a significant performance improvement was found when a scaling factor was applied to limb movement in the transverse plane. Also consistent with both the experimental hypothesis and previous work was the scaling value that produced the greatest performance increase - a scaling value of 1.5. This demonstrates that the initial observations regarding the depth judgement difficulties within the virtual environment, made in Section 1.3.4, and quantified and validated in Section 3.2, are both accurate and reliable. It validates the ecological validity of the experimental paradigm used throughout this thesis by reproducing effects found in the single-dimension tasks used throughout, and it also demonstrates the utility of the Stromohab research system in the investigation of motion-capture virtual-reality environments.

The quantification of this scaling factor has wide-reaching implications for designers of motion interfaced virtual environments, particularly virtual environments in which the minimisation of unwanted perceptual adjustment is desired. Virtual rehabilitation environments are an example of this; to restate an initial goal, it is highly desirable for a virtual environment to minimise additional cognitive load on a neurologically impaired patient, and furthermore, one of the aims of a virtual rehabilitation environment is to maximise the transference of useful functional motor skills to the physical world. Reducing and removing the deficit in the sagittal plane results in equal task performance in all three dimensions of movement, thereby eliminating the increased perceptual adjustment previously required during adaptation to a virtual environment. Reducing adjustment to the virtual environment then leads to the logical conclusion is that adjustment from the virtual environment, back to the physical world is also reduced. Thus, the transference of useful functional motor skills to the physical world is increased, and greater rehabilitative efficacy is achieved.

4.7 Concluding Discussion

This thesis presents a direct study of the egocentric scaling effect found in virtual environments, and presents a resolution to the impact of this effect on movement within a virtual environment. As described in Section 4.5 there are a number of experimental contexts in which somewhat similar observations have been noted, though not quantified. It is proposed that the well known rapid adaptation of the visuomotor and sensorimotor systems is a significant contributing factor to the absence of more widespread reporting of this effect. This proposal is further supported by Pelah and Barlow (1996), in which the authors show that the visuolocomotor system is itself adaptive, explaining the scarce observations of this effect in the growing area of locomotion simulator applications.

The adaptation of sensorimotor spatial relations can be observed after only a few minutes of exposure to spatial misalignment (Welch, 1978; Pelah and Barlow, 1996; Groen and Werkhoven, 1998). The classic examples of this effect were shown by Welch et al. (1974), and in more recent years Cunningham et al. (2001a) demonstrate temporal adaptation to misaligned cues. Visuomotor adaptation has even been shown to take place in eyes-closed conditions (Durgin and Pelah, 1999). It is therefore suggested that the two-factor perceptual compression observed and quantified in this thesis is likely present in many studies involving motion-controlled virtual environments, though is unlikely to have an effect on experimental outcomes due to rapid sensorimotor adaptation of virtual environment users.

Though this rapid adaptation may be viewed as advantageous in many virtual environment applications, and indeed in real environments, for which the mechanism evolved, when viewed through the lens of a rehabilitation system it immediately becomes a negative, and arguably dangerous effect.

A physical therapy patient recovering from neurological trauma such as stroke has by definition a motor skills deficit, and the primary goal of physical rehabilitation is to eliminate or minimise this deficit. For any form of physical therapy to be effective therefore, if motor skills can be reacquired they must be useful to the patient on their return to everyday life. While this is not generally an issue under conventional physiotherapy treatment protocols, motor skills acquired within a virtual environment may have no real-world application, and as such rehabilitative tasks must be carefully designed to maximise transference of function from the virtual to the real world.

Bias, or misrepresentation of movement within the virtual environment has the potential to impede rehabilitation in a number of ways. If neurological impairment is severe, adaptation to movement within an unfamiliar environment in which conflicting visual, haptic, and proprioceptive signals are experienced may be impeded by the additional cognitive load, to the point of impossibility. Further to this, should adaptation to a spatially misaligned environment take place, it follows that any skills learned within that environment will likely be most, if not only, applicable to that environment. Functional transference to the real world is again likely to be impeded, potentially to the point of impossibility.

To state this in more than the abstract - if a patient acquires obstacle avoidance skills from a virtual environment, through learning that lifting a foot by a given amount is enough to step onto a kerb for example, it is clear that falls are more likely if that learned height is not great enough. This not only negates any positive effect the virtual environment may have, but also produces the situation in which dangerous motor skills must be unlearned, which is in itself a task potentially more difficult than initial motor learning. Correct scaling of movement within the virtual environment, as quantified in Section 3.6.3 is therefore of great importance to the useful transference of functional motor skills to the physical world.

It is clear therefore that careful and correct design and implementation of virtual environments for rehabilitation is of paramount importance. The quantification and resolution of distorting factors such as the distance compression discussed in this thesis must be at the forefront of any virtual rehabilitative environment designer's mind. Additionally, in apparent contrast to the majority of virtual rehabilitation environments and systems described in current literature, the likely and potential cognitive abilities of the intended beneficiaries of such a system must be taken into account in order to be effective.

4.7.1 Future Work

The use of virtual reality in neurological trauma rehabilitation has only begun to filter through to clinical trial stage relatively recently. Although initial results are encouraging, there is a need for more in-depth, objective research in this field. Some attempts have been made been made to this end (Henderson et al., 2007), though it is difficult to avoid vague conclusions and questionable methods. Indeed, with regard to clinical methods, much of the work in this field appears to disregard research methods that include a control patient group for instance, instead preferring to focus on the effectiveness of a particular technological intervention.

While the definition, design, implementation and verification of the *Stromohab* project has occupied this thesis, it is clear that future studies aiming to explore the rehabilitative applications of the system should include objective clinical involvement and patient trials.

Use of VR in rehabilitation is maturing, and a number of encouraging studies and reviews have been completed in recent years, and with the current availability of accessible motion-tracking technologies, such as the Microsoft Kinect®, the influence of of both virtual-reality and motion-capture on rehabilitation looks set to increase further. While video-based markerless motion-tracking technology such as the Microsoft Kinect® present a much lower barrier to use, a number of technological challenges must be overcome before video-based tracking can be utilised as a direct replacement for marker-based tracking. This is currently an area of active research (Da Gama et al., 2012; Bó et al., 2011; Lange et al., 2011; Chang et al., 2011), and should those challenges be met, the modularity of the Stromohab system provides a smooth path for the adoption of new motion-capture technology.

It is exciting to note that (methodological) comparisons of the efficacy of VR rehabilitation and conventional therapy have favoured the VR approach (Sveistrup and McComas, 2003; Sveistrup, 2004a; Rose et al., 2005; Bryanton et al., 2006). In addition to this, it is hoped that this thesis will further contribute to the encouragement of the collaborative interdisciplinary approach required to further this field.

Glossary

Analysis Of Variance is a statistical test of whether or not the means of several groups are all equal, used in comparing three or more means for statistical significance.

Augmented Reality refers to a display in which simulated imagery, graphics, or symbology is superimposed on, or 'in' a video stream of an environment.

Center of Rotation the only point in a plane that remains unchanged under a rotation of the plane. In skeletal reconstruction from motion capture, these points are used to locate joints.

Hue, Saturation, and Brightness A device independant method of describing colour.

Inverse Kinematics a mathematical method to find the degrees of freedom of a system subject to kinematic constraints.

Virtual Reality a computer generated artificial or synthetic environment which stimulates any or all of a user's visual, aural and haptic perception.

References

J Accot and S Zhai. Refining Fitts' law models for bivariate pointing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 193–200. ACM, 2003.

F Affolter. Perceptual processes as prerequisites for complex human behaviour. *Disability & Rehabilitation*, 1981. URL `http://informahealthcare.com/doi/abs/10.3109/03790798109167107`.

J M Airey, J H Rohlf, and F P Brooks Jr. Towards image realism with interactive update rates in complex virtual building environments. *ACM SIGGRAPH Computer Graphics*, 24(2):41–50, 1990. ISSN 0897913515.

F Amirabdollahian, G T Gomes, and G R Johnson. The peg-in-hole: a VR-based haptic assessment for quantifying upper limb performance and skills. In *Rehabilitation Robotics, 2005. ICORR 2005. 9th International Conference on*, pages 422–425. IEEE, 2005.

C Armbrüster, M Wolter, T Kuhlen, W Spijkers, and B Fimm. Depth perception in virtual reality: distance estimations in peri-and extrapersonal space. *CyberPsychology & Behavior*, 11(1):9–15, 2008. ISSN 1094-9313.

S Aruin, Asha Sharma, Rob Larkins, and Gouri Chaudhuri. Knee position feedback: its effect on management of pelvic instability in a stroke patient. *Disability & Rehabilitation*, 22(15):690–692, 2000.

P Atallah, A Pelah, and A Wilkins. Visual stress symptoms from stereoscopic television. In *Proceedings of the International Conference on 3D Imaging.* IEEE Signal Processing Society, 2012.

S Attygalle, M Duff, T Rikakis, and J He. Low-cost, at-home assessment system with Wii Remote based motion capture. In *Virtual Rehabilitation, 2008*, pages 168–174. IEEE, 2008. ISBN 1424427002.

R T Azuma and Others. A survey of augmented reality. *Presence-Teleoperators and Virtual Environments*, 6(4):355–385, 1997.

A T Bahill, D Adler, and L Stark. Most naturally occurring human saccades have magnitudes of 15 degrees or less. *Investigative Ophthalmology & Visual Science*, 14(6):468–469, 1975. ISSN 1552-5783.

W Barfield and C Rosenberg. Judgments of azimuth and elevation as a function of monoscopic and binocular depth cues using a perspective display. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 37(1): 173–181, 1995. ISSN 0018-7208.

J V Basmajian. Clinical use of biofeedback in rehabilitation. *Psychosomatics*, 23(1):67–73, January 1982.

J V Basmajian. *Biofeedback: Principles and practice for clinicians.* Williams & Wilkins, 1989.

Peter J Bex, Isabelle Mareschal, and Steven C Dakin. Contrast gain control in natural scenes. *Journal of Vision*, 7(11), August 2007. doi: 10.1167/7.11.12. URL `http://www.journalofvision.org/content/7/11/12.abstract`.

A Billard and M Arbib. Mirror neurons and the neural basis for learning by imitation. *Stamenov and Gallese*, pages 343–352, 2002.

G Bingham and J L Romack. The rate of adaptation to displacement prisms remains constant despite acquisition of rapid calibration. *Journal of Experimental Psychology: Human Perception and Performance*, 25(5):1331, 1999. ISSN 1939-1277.

G P Bingham, A Bradley, M Bailey, and R Vinner. Accommodation, occlusion, and disparity matching are used to guide reaching: A comparison of actual versus virtual environments. *Journal of Experimental Psychology: Human Perception and Performance; Journal of Experimental Psychology: Human Perception and Performance*, 27(6):1314, 2001. ISSN 1939-1277.

A P L Bó, M Hayashibe, and P Poignet. Joint angle estimation in rehabilitation with inertial sensors and its integration with Kinect. In *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*, pages 3479–3483. IEEE, 2011. ISBN 1424441218.

Berta Bobath. *Adult hemiplegia: Evaluation and treatment.* 1990. URL `http://www.getcited.org/pub/102941477`.

R Boian, A Sharma, C Han, A Merians, G Burdea, S Adamovich, M Recce, M Tremaine, and H Poizner. Virtual reality-based post-stroke hand rehabilitation. *Studies in health technology and informatics*, pages 64–70, 2002. ISSN 0926-9630.

R F Boian, J E Deutsch, C S Lee, G C Burdea, and J Lewis. Haptic effects for virtual reality-based post-stroke rehabilitation. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2003. HAPTICS 2003. Proceedings. 11th Symposium on*, pages 247–253. IEEE, 2003. ISBN 0769518907.

R Boulic, P Fua, L Herda, M Silaghi, J S Monzani, L Nedel, and D Thalmann. An anatomic human body for motion capture. In *Proc. EMMSEC*, volume 98, 1998.

J W Boyse. Interference detection among solids and surfaces. *Communications of the ACM*, 22(1):3–9, 1979. ISSN 0001-0782.

L M P Braccialli, F Rebelo, and L M Pereira. Can Virtual Reality Methodologies Improve the Quality of Life of People with Disabilities? *Advances in Usability Evaluation*, 21:20, 2012. ISSN 1466560541.

R M Bracewell. Stroke: neuroplasticity and recent approaches to rehabilitation. *Journal of Neurology, Neurosurgery & Psychiatry*, 74(11):1465–1465, November 2003. ISSN 0022-3050. doi: 10.1136/jnnp.74.11.1465. URL `http://jnnp.bmj.com/content/74/11/1465.full`.

L Bradley, B B Hart, S Mandana, K Flowers, M Riches, and P Sanderson. Electromyographic biofeedback for gait training after stroke. *Clinical rehabilitation*, 12(1):11, 1998.

Michael Brainin and Wolf-Dieter Heiss. *Textbook of Stroke Medicine (Cambridge Medicine)*. Cambridge University Press, 2009. ISBN 0521518261. URL `http://www.amazon.co.uk/Textbook-Stroke-Medicine-Cambridge/dp/0521518261`.

D A Bray and B R Konsynski. Virtual worlds: multi-disciplinary research opportunities. *ACM SIGMIS Database*, 38(4):17–25, 2007. ISSN 0095-0033.

J Broeren, M Rydmark, and K S Sunnerhagen. Virtual reality and haptics as a training device for movement rehabilitation after stroke: a single-case study. *Archives of physical medicine and rehabilitation*, 85(8):1247–1250, 2004. ISSN 0003-9993.

S Brunnstrom. Associated reactions of the upper extremity in adult patients with hemiplegia; an approach to training. *The Physical therapy review*, 1956. URL `http://www.ncbi.nlm.nih.gov/pubmed/13322637`.

C Bryanton, J Bosse, M Brien, J McLean, A McCormick, and H Sveistrup. Feasibility, motivation, and selective motor control: virtual reality compared to conventional home exercise in children with cerebral palsy. *CyberPsychology & Behavior*, 9(2):123–128, 2006. ISSN 1094-9313.

G Burdea, V Popescu, V Hentz, and K Colbert. Virtual reality-based orthopedic telerehabilitation. *Rehabilitation Engineering, IEEE Transactions on*, 8(3): 430–432, 2000.

E Burns, S Razzaque, A T Panter, M C Whitton, M R McCallus, and F P Brooks Jr. The hand is more easily fooled than the eye: Users are more sensitive to visual interpenetration than to visual-proprioceptive discrepancy. *Presence: teleoperators & virtual environments*, 15(1):1–15, 2006.

I G Burnside, H S Tobias, and D Bursill. Electromyographic feedback in the remobilization of stroke patients: a controlled trial. *Arch Phys Med Rehabil*, 63(5):217–222, May 1982.

A Caramazza and M Coltheart. Cognitive neuropsychology twenty years on. *Cognitive Neuropsychology*, 23(1):3–12, 2006. ISSN 0264-3294.

JHS Carr, RB Shepherd, and B Roberta. A motor relearning programme for stroke. 1982. URL `http://pgim.srilankahealthrepository.org/handle/123456789/16389`.

M Carrozzo and F Lacquaniti. Virtual reality: a tutorial. *Electroencephalogr Clin Neurophysiol*, 109(1):1–9, February 1998.

Rachael Casey, Adar Pelah, Jonathan Cameron, and Joan Lasenby. Influence of step frequency on visual speed perception during locomotion. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*, APGV '10, page 160, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0248-7. doi: 10.1145/1836248.1836282. URL `http://doi.acm.org/10.1145/1836248.1836282`.

J H Challis. A procedure for determining rigid body transformation parameters. *J Biomech*, 28(6):733–737, June 1995.

A H S Chan and D K T So. Shape characteristics of useful field of view and visual search time. *Le travail humain*, 70(4):343–367, 2008. ISSN 2130561365.

Y J Chang, S F Chen, and J D Huang. A Kinect-based system for physical rehabilitation: A pilot study for young adults with motor disabilities. *Research in developmental disabilities*, 32(6):2566–2570, 2011. ISSN 0891-4222.

I W Charlton, P Tate, P Smyth, and L Roren. Repeatability of an optimised lower body model. *Gait Posture*, 20(2):213–221, October 2004. doi: 10.1016/j.gaitpost.2003.09.004. URL `http://dx.doi.org/10.1016/j.gaitpost.2003.09.004`.

Tom Cornsweet. *Visual perception*. Academic Press, New York, 1970. ISBN 9780121897505. URL `http://www.worldcat.org/title/visual-perception/oclc/106217`.

E Cotton and R Kinsman. Conductive education and adult hemiplegia. *Churchill Livingstone, Edinburgh*, 1983. URL `http://scholar.google.com/scholar?q="author:E+author:Cotton"Conductiveeducationandadulthemiplegia"\#1`.

Steven C Cramer and Jeff D Riley. Neuroplasticity and brain repair after stroke. *Current opinion in neurology*, 21(1):76–82, February 2008. ISSN 1350-7540. doi: 10.1097/WCO.0b013e3282f36cb6. URL `http://www.ncbi.nlm.nih.gov/pubmed/18180655`.

J J Crisco, X Chen, M M Panjabi, and S W Wolfe. Optimal marker placement for calculating the instantaneous center of rotation. *J Biomech*, 27(9):1183–1187, September 1994.

J H Crosbie, S Lennon, J R Basford, and S M McDonough. Virtual reality in stroke rehabilitation: still more virtual than real. *Disability & Rehabilitation*, 29(14):1139–1146, 2007. ISSN 0963-8288.

Douglas W Cunningham, Astros Chatziastros, Markus von der Heyde, and Heinrich H Bülthoff. Driving in the future: Temporal visuomotor adaptation and generalization. *Journal of Vision*, 1(2), November 2001a. doi: 10.1167/1.2.3. URL `http://www.journalofvision.org/content/1/2/3.abstract`.

DW Cunningham, VA Billock, and BH Tsou. Sensorimotor adaptation to violations of temporal contiguity. *Psychological Science*, 2001b. URL `http://pss.sagepub.com/content/12/6/532.short`.

A Da Gama, T Chaves, L Figueiredo, and V Teichrieb. Poster: Improving motor rehabilitation process through a natural interaction based system using Kinect sensor. In *3D User Interfaces (3DUI), 2012 IEEE Symposium on*, pages 145–146. IEEE, 2012. ISBN 1467312045.

Graeme Dandy. *Planning and Design of Engineering Systems*. Routledge, 2007. ISBN 0415405513. URL `http://books.google.com/books?id=6C4SRK-EEikC\&pgis=1`.

R B Davis. A gait analysis data collection and reduction technique. *Human Movement Science*, 10(5):575–587, 1991.

C M Dean and R B Shepherd. Task-related training improves performance of seated reaching tasks after stroke. A randomized controlled trial. *Stroke*, 28 (4):722–728, April 1997.

E Delahunt, K Monaghan, and B Caulfield. Ankle function during hopping in subjects with functional instability of the ankle joint. *Scand J Med Sci Sports*, 17(6):641–648, December 2007. doi: 10.1111/j.1600-0838.2006.00612.x. URL `http://dx.doi.org/10.1111/j.1600-0838.2006.00612.x`.

B Delaney. On the trail of the shadow woman: The mystery of motion capture. *IEEE Computer Graphics and Applications*, 18(5):14–19, 1998.

A R Den Otter, A C H Geurts, M De Haart, T Mulder, and J Duysens. Step characteristics during obstacle avoidance in hemiplegic stroke. *Experimental Brain Research*, 161(2):180–192, 2005. ISSN 0014-4819.

K Desouza. *Agile Information Systems*. Routledge, 2012. ISBN 0750682353.

J E Deutsch, J Latonio, G C Burdea, and R Boian. Post-stroke rehabilitation with the Rutgers Ankle System: a case study. *Presence: Teleoperators & Virtual Environments*, 10(4):416–430, 2001. ISSN 1054-7460.

R Dickstein, S Hocherman, T Pillar, and R Shaham. Stroke rehabilitation. Three exercise therapy approaches. *Physical Therapy*, 66(8):1233–1238, 1986. ISSN 0031-9023. URL `http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve\&db=PubMed\&dopt=Citation\&list\_uids=3737695`.

V Dietz and S J Harkema. Locomotor activity in spinal cord-injured persons. *Journal of Applied Physiology*, 96(5):1954–1960, 2004. ISSN 8750-7587.

H Dong, A Pelah, J Cameron, and J Lasenby. The perceptual influences on gait transition of step parameters and optic flow in virtual environment locomotion simulators. In *Proceedings of the 5th symposium on Applied perception in graphics and visualization*, pages 143–146. ACM, 2008. ISBN 1595939814.

Peter Dring. *VR Rehab Walker*. Unpublished, University of York, 2007.

Pamela W Duncan, William B Stason, Harold P Adams, Alan M Adelman, David N Alexander, Duane S Bishop, Leonard Diller, Nancy E Donaldson, Carl V Granger, and Audrey L Holland. *Post-Stroke Rehabilitation: Clinical Practice Guideline*. Aspen Publishers, 1995. ISBN 9780834208117.

F H Durgin and A Pelah. Visuomotor adaptation without vision? *Experimental Brain Research*, 127(1):12–18, July 1999.

E Dursun, N Hamamci, S Dnmez, O Tznalp, and A Cakci. Angular biofeedback device for sitting balance of stroke patients. *Stroke*, 27(8):1354–1357, August 1996.

Erbil Dursun, Nigar Dursun, and Duygu Alican. Effects of biofeedback treatment on gait in children with cerebral palsy. *Disabil Rehabil*, 26(2):116–120, January 2004.

J D Easton, J L Saver, G W Albers, M J Alberts, S Chaturvedi, E Feldmann, T S Hatsukami, R T Higashida, S C Johnston, and C S Kidwell. Definition and Evaluation of Transient Ischemic Attack A Scientific Statement for Healthcare Professionals From the American Heart Association/American Stroke Association Stroke Council; Council on Cardiovascular Surgery and Anesthesia; Council on Cardiov. *Stroke*, 40(6):2276–2293, 2009. ISSN 0039-2499.

Judi A Edmans, John R F Gladman, Sue Cobb, Alan Sunderland, Tony Pridmore, Dave Hilton, and Marion F Walker. Validity of a Virtual Environment for Stroke Rehabilitation. *Stroke*, 37(11):2770–2775, November 2006. doi: 10.1161/01.STR.0000245133.50935.65. URL `http://stroke.ahajournals.org/content/37/11/2770.abstract`.

D B Elliott, K Sanderson, and A Conkey. The reliability of the PelliRobson contrast sensitivity chart. *Ophthalmic and Physiological Optics*, 10(1):21–24, 1990. ISSN 1475-1313.

SR Ellis. Nature and origins of virtual environments: A bibliographical essay. *Computing Systems in Engineering*, 1991. URL `http://www.sciencedirect.com/science/article/pii/095605219190001L`.

C J Erkelens and H Collewijn. Motion perception during dichoptic viewing of moving random-dot stereograms. *Vision research*, 25(4):583–8, January 1985a. ISSN 0042-6989. URL `http://www.ncbi.nlm.nih.gov/pubmed/4060612`.

CJ Erkelens and H Collewijn. Eye movements and stereopsis during dichoptic viewing of moving random-dot stereograms. *Vision research*, 25:1689–1700, 1985b. URL `http://www.sciencedirect.com/science/article/pii/0042698985901415`.

Alberto Esquenazi and Nathaniel H Mayer. Instrumented assessment of muscle overactivity and spasticity with dynamic polyelectromyographic and motion analysis for treatment planning. *Am J Phys Med Rehabil*, 83(10 Suppl):S19—-S29, October 2004.

J A Fodor. *The modularity of mind.* MIT press, 1983. ISBN 0262560259.

J M Foley and R Held. Visually directed pointing as a function of target distance, direction, and available cues. *Attention, Perception, & Psychophysics*, 12(3): 263–268, 1972. ISSN 1943-3921.

J Freeman, S E Avons, D E Pearson, and W A IJsselsteijn. Effects of sensory information and prior experience on direct subjective ratings of presence. *Presence: Teleoperators & Virtual Environments*, 8(1):1–13, 1999. ISSN 1054-7460.

H Frenz, M Lappe, M Kolesnik, and T Bührmann. Estimation of travel distance from visual motion in virtual environments. *ACM Trans. Appl. Percept*, 3(4): 419–428, 2007.

C Gerber and P Matter. Biomechanical analysis of the knee after rupture of the anterior cruciate ligament and its primary repair. An instant-centre analysis of function. *J Bone Joint Surg Br*, 65(4):391–399, August 1983.

M Gérin-Lajoie, C L Richards, J Fung, and B J McFadyen. Characteristics of personal space during obstacle circumvention in physical and virtual environments. *Gait & posture*, 27(2):239–247, 2008. ISSN 0966-6362.

A. Gopnik, A.N. Meltzoff, and P.K. Kuhl. *The scientist in the crib: What early learning tells us about the mind.* William Morrow Paperbacks, 2000.

T Y Grechkin, T D Nguyen, J M Plumert, J F Cremer, and J K Kearney. How does presentation method and measurement protocol affect distance estimation in real and virtual environments? *ACM Transactions on Applied Perception (TAP)*, 7(4):26, 2010. ISSN 1544-3558.

W J Greenleaf and M A Tovar. Augmenting reality in rehabilitation medicine. *Artificial Intelligence in Medicine*, 6(4):289–299, 1994.

R L Gregory. Perceptual illusions and brain models. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 171(1024):279–296, 1968. ISSN 0080-4649.

Ben Grimes. Animating an Avatar for Stroke Rehabilitation. Technical report, University of York, 2008.

Joris Groen and Peter J Werkhoven. Visuomotor Adaptation to Virtual Hand Position in Interactive Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 7(5):429–446, October 1998. ISSN 1054-7460. doi: 10.1162/105474698565839. URL `http://dx.doi.org/10.1162/105474698565839`.

P Haggard, J Cockburn, J Cock, C Fordham, and D Wade. Interference between gait and cognitive tasks in a rehabilitating neurological population. *Journal of Neurology, Neurosurgery & Psychiatry*, 69(4):479–486, 2000. ISSN 1468-330X.

George B Hanna, Sami M Shimi, and Alfred Cuschieri. Randomised study of influence of two-dimensional versus three-dimensional imaging on performance of laparoscopic cholecystectomy. *The Lancet*, 351(9098):248–251, January 1998. ISSN 0140-6736. doi: http://dx.doi.org/10.1016/S0140-6736(97)08005-7. URL `http://www.sciencedirect.com/science/article/pii/S0140673697080057`.

R Held. Dissociation of visual functions by deprivation and rearrangement. *Psychological Research*, 31(4):338–348, 1968. ISSN 0340-0727.

R Held and N Durlach. Telepresence, time delay and adaptation. *Pictorial communication in virtual and real ...*, 1991. URL `http://books.google.co.uk/books?hl=en\&lr=\&id=8c6PClwCymcC\&oi=fnd\&pg=PA232\&dq=telepresence+time+delay+and+adaptation\&ots=ScKNEL6vAe\&sig=0ILfFZQ3G5J1rg-HhZdnbej-iA8`.

R Held and N Gottlieb. Technique for studying adaptation to disarranged hand-eye coordination. *Perceptual and Motor Skills*, 8(3):83–86, 1958. ISSN 0031-5125.

A Henderson, N Korner-Bitensky, and M Levin. Virtual reality in stroke rehabilitation: a systematic review of its effectiveness for upper limb motor recovery. *Topics in stroke rehabilitation*, 14(2):52–61, 2007. ISSN 1074-9357.

Suzana Herculano-Houzel and Roberto Lent. Isotropic fractionator: a simple, rapid method for the quantification of total cell and neuron numbers in the brain. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 25(10):2518–21, March 2005. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.4526-04.2005. URL `http://www.ncbi.nlm.nih.gov/pubmed/15758160`.

L Herda, P Fua, R Plnkers, R Boulic, and D Thalmann. Skeleton-based motion capture for robust reconstruction of human motion. *ca*, page 77, 2000.

L Herda, P Fua, R Plnkers, R Boulic, and D Thalmann. Using skeleton-based tracking to increase the reliability of optical motion capture. *Hum Mov Sci*, 20(3):313–341, June 2001.

S Hesse and D Uhlenbrock. A mechanized gait trainer for restoration of gait. *Journal of rehabilitation research and development*, 37(6):701–708, 2000. ISSN 0748-7711.

S Hesse, C Bertelt, and A Schaffrin. Restoration of gait in nonambulatory hemiparetic patients by treadmill training with partial body-weight support. *Archives of physical ...*, 1994. URL `http://ukpmc.ac.uk/abstract/MED/7944913`.

J Hidler, D Nichols, and M Pelliccio. Multicenter randomized clinical trial evaluating the effectiveness of the Lokomat in subacute stroke. *... and Neural Repair*, 2009. URL `http://nnr.sagepub.com/content/23/1/5.short`.

G Hirzinger, J Heindl, and K Landzettel. Predictive and knowledge-based telerobotic control concepts. In *Robotics and Automation, 1989. Proceedings., 1989 IEEE International Conference on*, pages 1768–1777. IEEE, 1989. ISBN 0818619384.

M K Holden. Virtual environments for motor rehabilitation: review. *Cyberpsychology & behavior*, 8(3):187–211, 2005. ISSN 1094-9313.

Scott Holmes. *Analysis of Asymmetry in Gait for Medical Diagnosis.* PhD thesis, University of York, 2007.

S Holzreiter. Calculation of the instantaneous centre of rotation for a rigid body. *J Biomech*, 24(7):643–647, 1991.

C I Howarth, Beggs W D A., and Bowden J M. The relationship between speed and accuracy of movement aimed at a target. *Acta Psychologica*, 35:207, 1971.

M Huo, J Verner, L Zhu, and M A Babar. Software quality and agile methods. In *Computer Software and Applications Conference, 2004. COMPSAC 2004. Proceedings of the 28th Annual International*, pages 520–525. IEEE, 2004. ISBN 0769522092.

W Ijsselsteijn, H de Ridder, R Hamberg, D Bouwhuis, and J Freeman. Perceived depth and the feeling of presence in 3DTV. *Displays*, 18(4):207–214, 1998. ISSN 0141-9382.

H Iwata. Locomotion interface for virtual environments. *Robotics research*, pages 275–282, 2000.

D Jack, R Boian, a S Merians, M Tremaine, G C Burdea, S V Adamovich, M Recce, and H Poizner. Virtual reality-enhanced stroke rehabilitation. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, 9(3): 308–318, October 2001. ISSN 1534-4320. doi: 10.1109/7333.948460. URL `http://www.ncbi.nlm.nih.gov/pubmed/11561668`.

Peter Jackson. The Hobbit: An Unexpected Journey, 2012.

G M Jensen, J Gwyer, K F Shepard, and L M Hack. Expert practice in physical therapy. *Physical therapy*, 80(1):28–43, 2000. ISSN 0031-9023.

S Jezernik and G Colombo. Robotic orthosis Lokomat: a rehabilitation and research tool. *…: Technology at the …*, 2003. URL `http://onlinelibrary.wiley.com/doi/10.1046/j.1525-1403.2003.03017.x/full`.

RE Kalman. Contributions to the theory of optimal control. *Bol. Soc. Mat. Mex*, 1960. URL `http://fotc.ath.cx/book/basar\_control/08.pdf`.

B Karan and M Vukobratovic. Calibration and accuracy of manipulation robot models–An overview. *Mechanism and Machine Theory*, 29(3):479–500, 1994.

N Katz, H Ring, Y Naveh, R Kizony, U Feintuch, and P L Weiss. Interactive virtual environment training for safe street crossing of right hemisphere stroke patients with unilateral spatial neglect. *Disabil Rehabil*, 27(20):1235–1243, October 2005. doi: 10.1080/09638280500076079. URL `http://dx.doi.org/10.1080/09638280500076079`.

Emily A Keshner. Virtual reality and physical rehabilitation: a new toy or a new research and rehabilitation tool? *J Neuroeng Rehabil*, 1(1):8, December 2004. doi: 10.1186/1743-0003-1-8. URL `http://dx.doi.org/10.1186/1743-0003-1-8`.

D L Kleinman, S Baron, and W H Levison. An optimal control model of human response part I: Theory and validation. *Automatica*, 6(3):357–369, 1970. ISSN 0005-1098.

JM Knapp and JM Loomis. Limited field of view of head-mounted displays is not the cause of distance underestimation in virtual environments. *Presence: Teleoperators & Virtual Environments*, 2004. URL `http://www.mitpressjournals.org/doi/abs/10.1162/1054746042545238`.

A Koenig and M Wellner. Virtual gait training for children with cerebral palsy using the Lokomat gait orthosis. *Studies in health …*, 2008. URL `http://www.ncbi.nlm.nih.gov/pubmed/18391287`.

R Koheil, P T Reg, and A R Mandel. Joint position biofeedback facilitation of physical therapy in gait training. *American journal of physical medicine*, 59 (6):288, 1980.

D Koller, G Klinker, E Rose, D Breen, R Whitaker, and M Tuceryan. Real-time vision-based camera tracking for augmented reality applications. In *Proceedings of the ACM symposium on Virtual reality software and technology*, pages 87–94. ACM, 1997. ISBN 089791953X.

B Kopp, A Kunkel, W Mhlnickel, K Villringer, E Taub, and H Flor. Plasticity in the motor system related to therapy-induced improvement of movement after stroke. *Neuroreport*, 10(4):807–810, March 1999.

L Kovar, J Schreiner, and M Gleicher. Footskate cleanup for motion capture editing. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 97–104. ACM, 2002. ISBN 1581135734.

C Krewer, F Müller, B Husemann, and S Heller. The influence of different Lokomat walking conditions on the energy expenditure of hemiparetic patients and healthy subjects. *Gait & posture*, 2007. URL `http://www.sciencedirect.com/science/article/pii/S0966636206003092`.

T Kuhlen and C Dohle. Virtual reality for physically disabled people. *Computers in biology and medicine*, 25(2):205–211, 1995.

V K B Kumar. Study on television viewing habits and their effects among high school children of Davanagere city. 2006.

M A Lafortune, P R Cavanagh, H J Sommer, and A Kalenak. Three-dimensional kinematics of the human knee during walking. *J Biomech*, 25(4):347–357, April 1992.

J C K Lai, J Woo, E Hui, and W M Chan. Telerehabilitationa new model for community-based stroke rehabilitation. *Journal of Telemedicine and Telecare*, 10(4):199–205, 2004. ISSN 1357-633X.

B Lange, C Y Chang, E Suma, B Newman, A S Rizzo, and M Bolas. Development and evaluation of low cost game-based balance rehabilitation tool using the Microsoft Kinect sensor. In *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*, pages 1831–1834. IEEE, 2011. ISBN 1424441218.

J Liepert, W H Miltner, H Bauder, M Sommer, C Dettmers, E Taub, and C Weiller. Motor cortex plasticity during constraint-induced movement therapy in stroke patients. *Neurosci Lett*, 250(1):5–8, June 1998.

A Liu, G Tharp, L French, S Lai, and L Stark. Some of what one needs to know about using head-mounted displays to improve teleoperator performance. *Robotics and Automation, IEEE Transactions on*, 9(5):638–648, 1993. ISSN 1042-296X.

D Lloyd-Jones, R Adams, M Carnethon, G De Simone, T B Ferguson, K Flegal, E Ford, K Furie, A Go, and K Greenlund. Heart Disease and Stroke Statistics2009 Update A Report From the American Heart Association Statistics Committee and Stroke Statistics Subcommittee. *Circulation*, 119(3):480–486, 2009. ISSN 0009-7322.

J P Lord and K Hall. Neuromuscular reeducation versus traditional programs for stroke rehabilitation. *Archives of Physical Medicine and Rehabilitation*, 67(2):88–91, 1986. URL `http://www.ncbi.nlm.nih.gov/pubmed/3954571`.

L Lunenburger and G Colombo. Biofeedback in gait training with the robotic orthosis Lokomat. *... in Medicine and ...*, 2004. URL `http://ieeexplore.ieee.org/xpls/abs\_all.jsp?arnumber=1404352`.

L Lunenburger and G Colombo. Clinical assessments performed during robotic rehabilitation by the gait training robot Lokomat. *..., 2005. ICORR 2005. ...*, 2005. URL `http://ieeexplore.ieee.org/xpls/abs\_all.jsp?arnumber=1501116`.

G Luo, F Vargas-Martin, and E Peli. The role of peripheral vision in saccade planning: Learning from people with tunnel vision. *Journal of vision*, 8(14), 2008. ISSN 1534-7362.

Francine Malouin, Carol L Richards, Bradford McFadyen, and Julien Doyon. [New perspectives of locomotor rehabilitation after stroke]. *Med Sci (Paris)*, 19(10):994–998, October 2003.

Maija Mäntyjärvi and Tarja Laitinen. Normal values for the Pelli-Robson contrast sensitivity test. *Journal of Cataract & Refractive Surgery*, 27(2): 261–266, February 2001. ISSN 0886-3350. doi: http://dx.doi.org/10.1016/S0886-3350(00)00562-9. URL `http://www.sciencedirect.com/science/article/pii/S0886335000005629`.

A Mayr, M Kofler, and E Quirbach. Prospective, blinded, randomized crossover study of gait rehabilitation in stroke patients using the Lokomat gait orthosis. *... and Neural Repair*, 2007a. URL `http://nnr.sagepub.com/content/21/4/307.short`.

Andreas Mayr, Markus Kofler, Ellen Quirbach, Heinz Matzak, Katrin Fröhlich, and Leopold Saltuari. Prospective, blinded, randomized crossover study of gait rehabilitation in stroke patients using the Lokomat gait orthosis. *Neurorehabilitation and neural repair*, 21(4):307–14, 2007b. ISSN 1545-9683. doi: 10.1177/1545968307300697. URL `http://www.ncbi.nlm.nih.gov/pubmed/17476001`.

R D McIntosh, K I McClements, H C Dijkerman, D Birchall, and A D Milner. Preserved obstacle avoidance during reaching in patients with left visual neglect. *Neuropsychologia*, 42(8):1107–1117, 2004. ISSN 0028-3932.

Alma S Merians, David Jack, Rares Boian, Marilyn Tremaine, Grigore C Burdea, Sergei V Adamovich, Michael Recce, and Howard Poizner. Virtual reality-augmented rehabilitation for patients following stroke. *Phys Ther*, 82(9):898–915, September 2002.

P Metherall, E A Dymond, and N Gravill. Posture control using electrical stimulation biofeedback: a pilot study. *J Med Eng Technol*, 20(2):53–59, 1996.

M Minsky. Telepresence. *Omni*, 2(9):45–52, 1980.

A Mirelman, P Bonato, and J E Deutsch. Effects of training with a robot-virtual reality system compared with a robot alone on the gait of individuals after stroke. *Stroke*, 40(1):169–174, 2009. ISSN 0039-2499.

B J Mohler, S H Creem-Regehr, and W B Thompson. The influence of feedback on egocentric distance judgments in real and virtual environments. In *Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, pages 9–14. ACM, 2006. ISBN 1595934294.

T Molet, R Boulic, and D Thalmann. A real time anatomical converter for human motion capture. In *Eurographics Workshop on Computer Animation and Simulation*, volume 96, pages 79–94, 1996.

R Montoya, P Dupui, B Pags, and P Bessou. Step-length biofeedback device for walk rehabilitation. *Med Biol Eng Comput*, 32(4):416–420, July 1994.

J Moreland and M A Thomson. Efficacy of electromyographic biofeedback compared with conventional physical therapy for upper-extremity function in patients following stroke: a research overview and meta-analysis. *Phys Ther*, 74 (6):534–537, June 1994.

J D Moreland, M A Thomson, and A R Fuoco. Electromyographic biofeedback to improve lower extremity function after stroke: a meta-analysis. *Arch Phys Med Rehabil*, 79(2):134–140, February 1998.

A M Moseley, A Stark, I D Cameron, and A Pollock. Treadmill training and body weight support for walking after stroke. *Stroke*, 34(12):3006, 2003. ISSN 0039-2499.

A Murata. Extending effective target width in Fitts' law to a two-dimensional pointing task. *International journal of human-computer interaction*, 11(2): 137–152, 1999.

E B Nash, G W Edwards, J A Thompson, and W Barfield. A review of presence and performance in virtual environments. *International Journal of human-computer Interaction*, 12(1):1–41, 2000. ISSN 1044-7318.

D S Nichols. Balance retraining after stroke using force platform biofeedback. *Phys Ther*, 77(5):553–558, May 1997.

B M Nigg and W Herzog. *Biomechanics of the musculo-skeletal system.* Wiley, 1999.

M M Panjabi, V K Goel, and S D Walter. Errors in kinematic parameters of a planar joint: guidelines for optimal experimental design. *J Biomech*, 15(7): 537–544, 1982a.

M M Panjabi, V K Goel, S D Walter, and S Schick. Errors in the center and angle of rotation of a joint: an experimental study. *J Biomech Eng*, 104(3): 232–237, August 1982b.

I S W Party. National clinical guideline for stroke, 2008.

R Pausch, T Crea, and M Conway. A literature survey for virtual environments-Military flight simulator visual systems and simulator sickness. *Presence: Teleoperators and Virtual Environments*, 1(3):344–363, 1992.

M B Pavithra. Study on the prevalence of refractive errors among school children of 7-15 years age group and spectacle wear compliance rate after intervention in the urban and rural field practice area of Dr. BR Ambedkar Medical College. 2011.

M D Peláez-Coca, P Sobrado-Calvo, and F Vargas-Martín. Optoelectronic aid for patients with severely restricted visual fields in daylight conditions. *Journal of Modern Optics*, 58(19-20):1871–1879, 2011. ISSN 0950-0340.

A Pelah. DISPLAY SYSTEM FOR VIRTUAL REALITY, August 2007.

A Pelah and H B Barlow. Visual illusion from running. *Nature*, 381(6580): 283, May 1996. doi: 10.1038/381283a0. URL `http://dx.doi.org/10.1038/381283a0`.

A Pelah, A Nanapragasam, J Cameron, and J Lasenby. Visualizations for locomotor learning with real time feedback in VR. In *Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization*, page 139. ACM, 2009.

E Peli. The visual effects of head-mounted display (HMD) are not distinguishable from those of desk-top computer display. *Vision Research*, 38(13):2053, 1998. ISSN 0042-6989.

T S Perry and G Zorpette. Oled TV arrives [2013 Tech To Watch]. *Spectrum, IEEE*, 50(1):46–47, 2013. ISSN 0018-9235.

A Pollock, G D Baer, P Langhorne, and V M Pomeroy. Physiotherapy treatment approaches for stroke. *Stroke: A Journal of Cerebral Circulation*, pages 519–520, 2008. URL `http://dx.doi.org/10.1161/STROKEAHA.107.492710`.

A J Prevo, S L Visser, and T W Vogelaar. Effect of EMG feedback on paretic muscles and abnormal co-contraction in the hemiplegic arm, compared with conventional physical therapy. *Scand J Rehabil Med*, 14(3):121–131, 1982.

D Rand, R Kizony, U Feintuch, N Katz, N Josman, A S Rizzo, and P L Weiss. Comparison of two VR platforms for rehabilitation: video capture versus HMD. *Presence: Teleoperators & Virtual Environments*, 14(2):147–160, 2005. ISSN 1054-7460.

K Rayner. Eye movements and attention in reading, scene perception, and visual search. *The quarterly journal of experimental psychology*, 62(8):1457–1506, 2009. ISSN 1747-0218.

Adam R Richardson and David Waller. The effect of feedback training on distance estimation in virtual environments. *Applied Cognitive Psychology*, 19(8):1089–1108, 2005. ISSN 08884080. doi: 10.1002/acp.1140. URL `http://doi.wiley.com/10.1002/acp.1140`.

B E Riecke, J Schulte-Pelkum, M N Avraamides, M V D Heyde, and H H Bülthoff. Cognitive factors can influence self-motion perception (vection) in virtual reality. *ACM Transactions on Applied Perception (TAP)*, 3(3):194–216, 2006. ISSN 1544-3558.

R Riener and L Lünenburger. Locomotor training in subjects with sensorimotor deficits: an overview of the robotic gait orthosis lokomat. *Journal of Healthcare . . .*, 2010. URL `http://multi-science.metapress.com/index/T2Q26427T72J6MQ6.pdf`.

J J Rieser, D H Ashmead, C R Talor, and G A Youngquist. Visual perception and the guidance of locomotion without vision to previously seen targets. *Perception*, 19(5):675–689, 1990.

Giuseppe Riva, Fabrizia Mantovani, and Andrea Gaggioli. Presence and rehabilitation: toward second-generation virtual reality applications in neuropsychology. *Journal of Neuroengineering and Rehabilitation*, 1(1):1–11, 2004. doi: 10.1186/1743-0003-1-9.

Rood. Neurologic reactions: a basis for physical therapy, 1954. URL `http://journals.cambridge.org/action//displayFulltext?fromPage=online\&type=6\&fid=S0959259800004135\&aid=2679972\&next=true\&jid=RCG\&volumeId=5\&issueId=02\&next=Y`.

E Ros, J Díaz, S Mota, F Vargas-Martin, and M Peláez-Coca. Real time image processing on a portable aid device for low vision patients. *Reconfigurable Computing: Architectures and Applications*, pages 158–163, 2006.

F D Rose, E A Attree, and D A Johnson. Virtual reality: an assistive technology in neurological rehabilitation. *Current Opinion in Neurology*, 9(6):461, 1996.

F David Rose, Barbara M Brooks, and Albert A Rizzo. Virtual reality in brain damage rehabilitation: review. *Cyberpsychology and Behavior*, 8(3): 241–262; discussion 263–271, 2005. URL `http://www.ncbi.nlm.nih.gov/pubmed/15971974`.

L B Rosenberg. The effect of interocular distance upon operator performance using stereoscopic displays to perform virtual depth tasks. In *Virtual Reality Annual International Symposium, 1993., 1993 IEEE*, pages 27–32. IEEE, 1993. ISBN 0780313631.

C S Sahm, S H Creem-Regehr, W B Thompson, and P Willemsen. Throwing versus walking as indicators of distance perception in similar real and virtual environments. *ACM Transactions on Applied Perception (TAP)*, 2(1):35–45, 2005. ISSN 1544-3558.

M V Sanchez-Vives and M Slater. From presence to consciousness through virtual reality. *Nature Reviews Neuroscience*, 6(4):332–339, 2005. ISSN 1471-003X.

S Sandeep and L Mindy. Viewing medium affects arm motor performance in 3D virtual environments. *Journal of NeuroEngineering and Rehabilitation*, 8. ISSN 1743-0003.

Richard A Schmidt and Craig A Wrisberg. *Motor Learning and Performance w/Web Study Guide - 4th Edition: A Situation-Based Learning Approach*. Human Kinetics, 4 edition, 2007. ISBN 9780736069649.

M J Schuemie, P Van Der Straaten, M Krijn, and C.A.P.G. Van Der Mast. Research on presence in virtual reality: A survey. *CyberPsychology & Behavior*, 4(2):183–201, 2001. ISSN 1094-9313.

M T Schultheis, J Himelstein, and A A Rizzo. Virtual reality and neuropsychology: upgrading the current tools. *The Journal of head trauma rehabilitation*, 17(5):378, 2002. ISSN 0885-9701.

MT Schultheis and AA Rizzo. The application of virtual reality technology in rehabilitation. *Rehabilitation Psychology*, 2001. URL `http://psycnet.apa.org/journals/rep/46/3/296/`.

P Servos. Distance estimation in the visual and visuomotor systems. *Experimental Brain Research*, 130(1):35–47, 2000. ISSN 0014-4819.

Satoshi Shioiri, Masanori Yoshizawa, Mistuharu Ogiya, Kazumichi Matsumiya, and Hirohisa Yaguchi. Low-level motion analysis of color and luminance for perception of 2D and 3D motion. *Journal of Vision*, 12(6), June 2012. doi: 10.1167/12.6.33. URL `http://www.journalofvision.org/content/12/6/33.abstract`.

M Slater, M Usoh, and Y Chrysanthou. The influence of dynamic shadows on presence in immersive virtual environments. In *Selected papers of the Eurographics workshops on Virtual environments*, volume 95, pages 8–21. Citeseer, 1995.

M Sliger. Agile Projects in the Waterfall Enterprise. *BETTER SOFTWARE*, 2006.

Suzanne C. O'Connell Smeltzer, Brenda G. Bare, Janice L. Hinkle, and Kerry H. Cheever. *Brunner and Suddarth's Textbook of Medical Surgical Nursing*. Lippincott Williams & Wilkins, 2010. ISBN 0781785898. URL `http://books.google.com/books?id=SmtjSD1x688C\&pgis=1`.

J J Spiegelman and S L Woo. A rigid-body method for finding centers of rotation and angular displacements of planar joint motion. *J Biomech*, 20(7):715–721, 1987.

T Starner, S Mann, B Rhodes, J Levine, J Healey, D Kirsch, R W Picard, and A Pentland. Augmented reality through wearable computing. *Presence: Teleoperators and Virtual Environments*, 6(4):386–398, 1997. ISSN 1054-7460.

M Steinisch, M G Tana, and S Comani. A passive robotic device for VR-augmented upper limb rehabilitation in stroke patients. 2012. ISSN 1862-278X.

PH Stern and F McDowell. Effects of facilitation exercise techniques in stroke rehabilitation. *. . . and rehabilitation*, 1970. URL `http://ukpmc.ac.uk/abstract/MED/5475721`.

J Steuer. Defining virtual reality: Dimensions determining telepresence. *Journal of communication*, 42(4):73–93, 2006. ISSN 1460-2466.

R T Stone, K P Watts, and B E Rosenquist. Evaluation of 3D Television: Impact on Depth Perception. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 56, pages 733–737. SAGE Publications, 2012. ISBN 1071-1813.

K Stuart, K William, and J BETTY. Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT. *Ergonomics*, 21(8):601–613, 1978.

M Sumrell. From waterfall to agile-how does a qa team transition? In *Agile Conference (AGILE), 2007*, pages 291–295. IEEE, 2007. ISBN 0769528724.

K Sureshchandra and J Shrinivasavadhani. Moving from waterfall to agile. In *Agile, 2008. AGILE'08. Conference*, pages 97–101. IEEE, 2008. ISBN 0769533213.

H Sveistrup. Review: Motor rehabilitation using virtual reality. *Journal of NeuroEngineering and Rehabilitation*, 1:10, 2004a.

H Sveistrup and J McComas. Experimental studies of virtual reality-delivered compared to conventional exercise programs for rehabilitation. *CyberPsychology & . . .*, 2003. URL `http://online.liebertpub.com/doi/abs/10.1089/109493103322011524`.

H Sveistrup, M Thornton, C Bryanton, J McComas, S Marshall, H Finestone, A McCormick, J McLean, M Brien, Y Lajoie, and E Bisson. Outcomes of intervention programs using flatscreen virtual reality. *Conf Proc IEEE Eng Med Biol Soc*, 7:4856–4858, 2004. doi: 10.1109/IEMBS.2004.1404343. URL `http://dx.doi.org/10.1109/IEMBS.2004.1404343`.

Heidi Sveistrup. Motor rehabilitation using virtual reality. *J Neuroeng Rehabil*, 1(1):10, December 2004b. doi: 10.1186/1743-0003-1-10.

Heidi Sveistrup, Joan McComas, Marianne Thornton, Shawn Marshall, Hillel Finestone, Anna McCormick, Kevin Babulic, and Alain Mayhew. Experimental studies of virtual reality-delivered compared to conventional exercise programs for rehabilitation. *Cyberpsychol Behav*, 6(3):245–249, June 2003. doi: 10.1089/109493103322011524. URL `http://dx.doi.org/10.1089/109493103322011524`.

M J Tarr and W H Warren. Virtual reality in behavioral neuroscience and beyond. *Nature Neuroscience*, 5:1089–1092, 2002.

E Taub, N E Miller, T A Novack, E W Cook, W C Fleming, C S Nepomuceno, J S Connell, and J E Crago. Technique to improve chronic motor deficit after stroke. *Arch Phys Med Rehabil*, 74(4):347–354, April 1993. URL `http://ukpmc.ac.uk/abstract/MED/8466415`.

L F Teixeira-Salmela, S J Olney, S Nadeau, and B Brouwer. Muscle strengthening and physical conditioning to reduce impairment and disability in chronic stroke survivors. *Arch Phys Med Rehabil*, 80(10):1211–1218, October 1999.

M Teschner, S Kimmerle, B Heidelberger, G Zachmann, L Raghupathi, A Fuhrmann, M P Cani, F Faure, N MagnenatThalmann, and W Strasser. Collision detection for deformable objects. In *Computer Graphics Forum*, volume 24, pages 61–81. Wiley Online Library, 2005. ISBN 1467-8659.

WB Thompson and P Willemsen. Does the quality of the computer graphics matter when judging distances in visually immersive environments? *. . . & Virtual Environments*, 2004. URL `http://www.mitpressjournals.org/doi/abs/10.1162/1054746042545292`.

A Tognetti, F Lorussi, R Bartalesi, S Quaglini, M Tesconi, G Zupone, and D De Rossi. Wearable kinesthetic system for capturing and classifying upper limb gesture in post-stroke rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 2(1):8, 2005. ISSN 1743-0003.

B J Unger, A Nicolaidis, P J Berkelman, A Thompson, R L Klatzky, and R L Hollis. Comparison of 3-D haptic peg-in-hole tasks in real and virtual environments. In *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, volume 3, pages 1751–1756. IEEE, 2001.

B J Unger, A Nicolaidis, P J Berkelman, A Thompson, S Lederman, R L Klatzky, and R L Hollis. Virtual peg-in-hole performance using a 6-dof magnetic levitation haptic device: Comparison with real forces and with visual guidance alone. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2002. HAPTICS 2002. Proceedings. 10th Symposium on*, pages 263–270. IEEE, 2002.

K I Ustinova, W A Leonard, N D Cassavaugh, and C D Ingersoll. Development of a 3D immersive videogame to improve arm-postural coordination in patients with TBI. *Journal of neuroengineering and rehabilitation*, 8(1):61, 2011. ISSN 1743-0003.

A Väljamäe. Self-motion and presence in the perceptual optimization of a multisensory virtual reality environment, 2005.

RPS van Peppen. Towards evidence-based physiotherapy for patients with stroke. 2008. URL `http://igitur-archive.library.uu.nl/dissertations/2008-0530-200239/UUindex.html`.

G Vaněkček Jr. Backface culling applied to collision detection of polyhedra. *The Journal of Visualization and Computer Animation*, 5(1):55–63, 1994. ISSN 1099-1778.

F Vargas-Martin. Augmented-view for restricted visual field: multiple device implementations. *Optometry & Vision Science*, 79(11):715, 2002. ISSN 1040-5488.

F Vargas-Martín, M D Peláez-Coca, E Ros, J Diaz, and S Mota. A generic real-time video processing unit for low vision. In *International Congress Series*, volume 1282, pages 1075–1079. Elsevier, 2005. ISBN 0531-5131.

F E Veldpaus, H J Woltring, and L Dortmans. A least-squares algorithm for the equiform transformation from spatial marker co-ordinates. *Journal of Biomechanics*, 21(1):45–54, 1988.

J P Wann and J D Turnbull. Motor skill learning in cerebral palsy: movement, action and computer-enhanced therapy. *Bailli{è}re's clinical neurology*, 2(1): 15, 1993.

Charles P. Warlow, Jan van Gijn, Martin S. Dennis, Joanna M. Wardlaw, John M. Bamford, Graeme J. Hankey, Peter A. G. Sandercock, Gabriel Rinkel, Peter Langhorne, Cathie Sudlow, and Peter Rothwell. *Stroke: Practical Management.* John Wiley & Sons, 2011. ISBN 1444359258. URL `http://books.google.com/books?hl=en\&lr=\&id=Br5DUstIb9QC\&pgis=1`.

P L Weiss, D Rand, N Katz, and R Kizony. Video capture virtual reality as a flexible and effective rehabilitation tool. *Journal of neuroengineering and rehabilitation*, 1(1):12, 2004a. ISSN 1743-0003.

Patrice L Weiss, Debbie Rand, Noomi Katz, and Rachel Kizony. Video capture virtual reality as a flexible and effective rehabilitation tool. *Journal of neuroengineering and rehabilitation*, 1(1):12, December 2004b. ISSN 1743-0003. doi: 10.1186/1743-0003-1-12. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=546410\&tool=pmcentrez\&rendertype=abstract`.

G Welch and E Foxlin. Motion tracking survey. *IEEE Computer Graphics and Applications*, pages 24–38, 2002.

R B Welch. *Perceptual modification: Adapting to altered sensory environments.* Academic Press New York, 1978. ISBN 0127418504.

R B Welch, C S Choe, and D R Heinrich. Evidence for a three-component model of prism adaptation. *Journal of Experimental Psychology*, 103(4):700, 1974. ISSN 0022-1015.

R B Welch, T T Blackmon, A Liu, B A Mellers, and L W Stark. The effects of pictorial realism, delay of visual feedback, and observer interactivity on the subjective sense of presence. *Presence: Teleoperators and Virtual Environments*, 5(3):263–273, 1996. ISSN 1054-7460.

D West, T Grant, M Gerush, and D DSilva. Agile development: Mainstream adoption has changed agility. *Forrester Research*, 2010.

B L White, P Castle, and R Held. Observations on the development of visually-directed reaching. *Child Development*, pages 349–364, 1964. ISSN 0009-3920.

P N Wilson, N Foreman, and M Tlauka. Transfer of spatial information from a virtual to a real environment in physically disabled children. *Disabil Rehabil*, 18(12):633–637, December 1996.

C J Winstein, E R Gardner, D R McNeal, P S Barto, and D E Nicholson. Standing balance training: effect on balance and locomotion in hemiparetic adults. *Arch Phys Med Rehabil*, 70(10):755–762, October 1989.

BG Witmer and PB Kline. Judging perceived and traversed distance in virtual environments. *Presence*, 1998. URL `http://www.mitpressjournals.org/doi/abs/10.1162/105474698565640`.

Bob G Witmer and Wallace J Sadowski. Nonvisually Guided Locomotion to a Previously Viewed Target in Real and Virtual Environments. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 40(3):478–488, September 1998. doi: 10.1518/001872098779591340. URL `http://hfs.sagepub.com/content/40/3/478.abstract`.

S L Wolf. Electromyographic biofeedback applications to stroke patients. A critical review. *Phys Ther*, 63(9):1448–1459, September 1983.

S L Wolf and S A Binder-MacLeod. Electromyographic biofeedback applications to the hemiplegic patient. Changes in lower extremity neuromuscular and functional status. *Phys Ther*, 63(9):1404–1413, September 1983.

G Wu. Real-time feedback of body center of gravity for postural training of elderly patients with peripheral neuropathy. *IEEE Trans Rehabil Eng*, 5(4): 399–402, December 1997.

J Xu, D Liu, and G Fang. An efficient method for collision detection and distance queries in a robotic bridge maintenance system. *Robotic Welding, Intelligence and Automation*, pages 71–82, 2007.

H Yano, K Kasai, H Saitou, and H Iwata. Development of a gait rehabilitation system using a locomotion interface. *The Journal of Visualization and Computer Animation*, 14(5):243–252, 2003. ISSN 1099-1778.

S H You, S H Jang, Y H Kim, M Hallett, S H Ahn, Y H Kwon, J H Kim, and M Y Lee. Virtual realityinduced cortical reorganization and associated locomotor recovery in chronic stroke an experimenter-blind randomized study. *Stroke*, 36(6):1166–1171, 2005. ISSN 0039-2499.

S Zhai. On the validity of throughput as a characteristic of computer input. *IBM Research Report RJ*, 10253, 2002.

S Zhang, K H Leo, B Y Tan, and R Q F Tham. A hybrid human motion tracking system for virtual rehabilitation. In *Industrial Electronics and Applications (ICIEA), 2011 6th IEEE Conference on*, pages 1993–1998. IEEE, 2011. ISBN 1424487544.

H Zhou and H Hu. A survey-human movement tracking and stroke rehabilitation. *Department report CSM-420, University of Essex, England, ISSN*, pages 1744–8050, 2004.

H Zhou and H Hu. Human motion tracking for rehabilitation–A survey. *Biomedical Signal Processing and Control*, 3(1):1–18, 2008.

V B Zordan and N C Van Der Horst. Mapping optical motion capture data to skeletal motion using a physical model. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 245–250. Eurographics Association, 2003.

V B Zordan, A Majkowska, B Chiu, and M Fast. Dynamic response for motion capture animation. *ACM Transactions on Graphics (TOG)*, 24(3):697–701, 2005. ISSN 0730-0301.