Reducing jerk for bimanual control with virtual reality

Sean Chase Mandrake Hill

A thesis submitted in partial fulfilment of the requirement of the
University of Salford for the degree of Doctor of Philosophy
undertaken within the school of Computing, Science and Engineering

January 2019

Table of Contents

Abstract

This work explores interaction in virtual environments where bimanual operations lacking kinematic constraints, such as a pair of tracked controllers held in free space, are mapped to a target object that obeys a constraint, such as a pair of grippers grasping a single object.

Visual and haptic feedback systems are developed and compared as solutions for minimising Jerk during motion. The outcome of the study is that visual feedback is shown to be superior to haptic feedback for use in free-air bimanual systems where maximising smoothness of motion is the objective.

The primary application is for improved teleoperation of a new generation of ordnance disposal robots which utilise dual grippers intended for bimanual operation. The work is undertaken in a generalised manner however so that it may be applicable for other uses such as medical systems or nuclear disassembly.

The systems developed herein initially utilise a custom framework designed for a fourteen screen, fully tracked immersive environment comparable to an enhanced CAVE, known locally as Octave. The final study subsequently utilises the Unreal Engine 4 framework for the HTC Vive HMD system. All software developed is provided as a core part of this work.

Declaration

I declare that this thesis has been composed solely by myself and that it has not been submitted, in whole or in part, in any previous application. Except where stated otherwise by reference or acknowledgement, the work presented is entirely my own.

Sean Chase Mandrake Hill

January 2019

1. Introduction

The majority of research done to date regarding bimanual interaction focusses upon mapping a consistent haptic resistance to the user so that the kinematic chain is echoed directly. This includes such systems as bimanual surgical systems. In these instances it may be considered that the aim of a system is to be transparent to the user, creating minimal impact, so that the user may leverage their training and skills to maximise performance.

The work here deviates from this and explores the idea that the system itself can be leveraged to enhance aspects of operator performance if it takes advantage of the unique features of a virtual workspace where input to output mapping can be arbitrary.

A new framework for research is developed and two pilot studies are undertaken. These studies are primarily confirmation studies. The first study explores user preference for bimanual interaction and confirms that users tend to prefer bimanual interaction more often with repeated use. The second study extends the first with the addition of a physics model and feedback indicators. The results of this confirm that visual feedback improves user precision if well implemented.

The majority of prior works explored focus upon either precision or speed as metrics. In a number of these works users are asked to be both fast and precise despite these factors being oppositional. The pilot studies avoid this ambiguity by focussing only upon precision but this is not a rigorous solution since users may still limit their attempt time to what they personally consider reasonable. As an alternative approach the core work here utilises the jerk metric instead, which is growing in popularity amongst medical systems evaluation but has yet to be adopted comprehensively elsewhere.

The jerk metric evaluates smoothness of motion to offer a usefully composite measurement whereby a slow-but-smooth interaction might be ranked more desirable that a swift-but-imprecise one, or indeed vice-versa, depending upon the motion of the user over time by comparison to a mathematical ideal, rather than to an arbitrary value or empiric ratio.

The core study therefore develops a system to map bimanual unconstrained motion to a real-world constrained object whilst comparing different forms of supporting feedback to optimise operator interaction with the aim of minimising jerk.

1.1 Hypothesis

This study investigates methods of reducing jerk under bimanual operation in free space.

The study proposes that the use of feedback indicators will have a significant and positive impact upon the smoothness of user performance. The primary hypothesis in its null and alternative forms is thus given as:

Ha_0 Jerk levels will be improved through the use of indicators

Ha_1 Jerk levels will be inferior or equal without indicators

This study proposes as its secondary hypothesis that visual feedback will have a more positive impact upon performance than haptic feedback. The secondary hypothesis in its null and alternative forms can thus be given as:

Hb_0 Jerk levels will be improved with visual indicators

Hb_1 Jerk levels will be reduced with haptic indicators

1.2 Application

The primary application for this work is teleoperation of the current standard of Explosive Ordnance Disposal platforms. The primary method of ordnance disposal is to move the device to a safe location where it can be detonated in a controlled manner. In order to do this it is essential that the operator has clear information regarding the jerk which the platform is subjecting the device to.

This work has secondary relevance to medical platforms also. The majority of these platforms focus upon using proportional haptic feedback to resist force and utilise static input systems fixed to a surface in order to enable this. The concept of utilising methods other than force resistance to represent feedback information may be highly applicable.

1.3 Novelty

This is a new study into free-space bimanual interaction regarding a single objects that focusses upon reducing jerk by using different forms of user feedback.

Few studies investigate bimanual control in free-space and those that do focus largely upon usability studies, preferring user feedback and qualitative analysis.

Most bimanual studies either analyse the hands separately to find better ways of utilising them in conjunction or else assume a hard kinematic constraint. This study explores the use of an unconstrained pair of hands mapped to a single target object.

Studies concerned with minimising jerk predominantly assume the desirability of haptic feedback whereas this study explores an alternative to traditional haptic feedback.

1.4 Contributions

The primary contribution of this study is a better understanding of the impact indicators and feedback have upon users with regard to reducing jerk. The study demonstrates that visual indicators can be superior to haptic indicators whilst confirming that both are superior to none.

The study also contributes a new framework design for evaluating free-space motion that is re-usable as a benchmark and a work of software that is reusable and easily modified for further experiments. The software accessible through the links in Appendix I contains a new experiment that is ready to run without the need for any modification or changes and will immediately extend the work covered herein to include exploration into hand placement and movement during the jerk analysis process.

2. Literature Review

2.1 The state of the art

Much research for use within Immersive Virtual Environments tends to focus designing technology rather than addressing specific use cases. It is common to assume that any effective technology will naturally find a space within an ecosystem, or that if the focus is upon improving a technology then, once sufficiently advanced, the inherent nature of the technology itself will announce self-evident application and value.

Doug Bowman, perhaps the most well-published researcher in the field, argues that the slowing of research after the 1990’s boom of innovation into 3D user interfaces is evidence that the design space of possibilities has now been largely covered (Bowman et al., 2006). He further suggests that the failure to adopt any particular system as a de-facto standard is evidence that the process of refinement for usability, by contrast, has so far been unsuccessful. Or in other words, a plethora of novel solutions exists but none have been pursued and refined to the point they are clearly superior to the others. Bowman attempts to summarise and analyse the general state research thus far and makes the case that the shortcomings of 3DUI, as it currently stands, will only be addressed through analysis of what is needed and a change in direction accordingly. He believes that if the enhancement and evolution of existing systems could be met simply by generating new ideas then we would have seen ubiquitous standards and universal systems adopted and accepted by now.

The further observation is made that perhaps over-generalisation of research has led to implementations that fail to address the specific needs of any domain, task or application at all. The outcome of this being that regardless of whatever advantage or improvement has been made, shortcomings exist within each research work that render them flawed from a usability perspective. Bowman creates an advanced graduate seminar class with ten students, in teams of two, allocated projects of a specific nature with the goal of improving specific 3DUI areas rather than of creating new technologies. The outcomes evidences more progress than is considered typical (Lucas et al., 2005; Polys et al., 2007).

Bowman has gone on to lay down simple guidelines for future research a number of times. He is insistent that researchers must focus upon usability, not technology:

> *"We stress that researchers must continue to perform evaluations of usability and performance, and that further guidelines for the development of effective 3-D interfaces are necessary. Such research has a direct impact on the acceptance of VEs as real-world tools, and on the levels of user satisfaction and productivity in VE applications."*

(Bowman et al., 2001)

Another factor that Bowman's commentary does not address is that access to a research-grade Virtual Environment (VE) technologies can be difficult to secure. Such a system requires ongoing support and resources to ensure that what is learned from year to year, especially with regard to in-house configuration and practical matters, need not be needlessly repeated. The technology itself should be a transparent, calibrated tool ready for useful research use, but in reality, preparing the VE system for research may consume the bulk of the researcher's time.

In many of the papers cited throughout this work it can be observed that unreasonable effort is often expended developing software and hardware to provide the essential foundations to enable research. Bowman's observation that a significant number of research papers focus excessively upon the development of new technology rather than refinement of knowledge is perhaps overly harsh and a failure to acknowledge the significant difficulties associated with readying such technology. The outlay necessary to maintain a fully featured immersive environment is significant[1] and requires continual investment to ensure the facilities are research-ready at all times. VR researchers must be prepared to put in hours far above and beyond anything that could be considered typical.

[1] When this work began a complicated and expensive cluster system with projectors, tracking and bespoke software was the only reasonable option. By completion a tracked head mounted system like the HTC Vive or Oculus Rift, with a variety of software options, is available for around one-thousandth the cost.

Bowman expands upon his guidelines previously presented at the first IEEE Symposium on 3D User Interfaces (Bowman et al., 2006) and in his earlier publications (Bowman et al., 2005). In particular Bowman laid down eight ideological guidelines towards attaining high performance 3D interaction (Bowman et al., 2008).

1. Floating objects should exist only as an exception.
2. Solid objects should not interpenetrate.
3. Only visible objects should interact.
4. Perspective and occlusion will provide the strongest depth cues.
5. People focus on the object, not the cursor.
6. 3D rotations are not always necessary: Objects may have inherent constraints.
7. 2D tasks are cognitively simpler than 3D and should be used where possible.
8. 2D input devices have been proven to be more precise.

Despite Bowman's extensive work and a wealth of supporting these guidelines are rarely adhered to. Bowman acknowledges the continuing immaturity of immersive systems however and implies that meeting these guidelines continues to be a significant challenge. Domain Specific Design is proposed as a means to address the balance of overly specific versus overly general research (Chen and Bowman, 2009).

The SPIDAR series of interfaces (Liu et al., 2014; Murayama et al., 2004) provide an example of highly focussed research. The multi-finger haptic interface developed here claims no domain applicability and offers no general insights but instead addresses only its own exactingly specific nature. The SPIDAR series is far from unique in this and in a thorough survey of more than thirty techniques (Argelaguet and Andujar, 2013) it was of key concern that most of these systems examined were evaluated in isolation. There were no de-facto benchmarks, no alternative datasets for comparison, and both hardware and software varied with every work.

3D touch, miniature wearable controller for immersive systems, is another example of isolated specificity (Nguyen and Banic, 2015). Whilst it certainly meets the requirements of valid research in that it demonstrates novelty and offers a new concept, it is difficult to see how this is of any merit in practice. Miniaturisation is nothing new, nor are mouse-type

OPS trackers or gyro-type IMU motion sensors. There is justifiable pride in having created such complex hardware, but without any means to recreate and build upon the work the value seems questionable.

Cho introduces another novel piece of bimanual hardware and follows the established pattern of empirically comparing it to systems that have gone before (Cho and Wartell, 2015). The work is rigorous and well-validated but offers only incremental improvements to systems that have thus-far failed to be adopted. The conclusion shows simply, without insight, that this new hardware offers marginally improved performance in some areas and marginally worse in others. It is difficult to see how such a work can be built upon.

An older work of particular interest returns to basic principles to develop a versatile, immersive, bimanual platform for experimental research purposes, to grab move and manipulate objects (Boussemart et al., 2004). This work is highly comprehensive but hindered by the equipment of the time. The author notes how consumer graphics cards are just beginning to open doors to new research that were unavailable until recently whilst lamenting the low rendering performance and limited fidelity of tracking of the time. What the author lacks in technology is made up for through thorough application of research principles however and thus this older work remains instructional in its execution.

The Reality-based User Interface System RUIS (Takala, 2014), a recent VR toolkit, is driven by the need for “a VR toolkit with a *low barrier of entry*”. It is made clear that the research department of Aalto University had become highly frustrated by system and software incompatibilities and overly complex development problems halting and slowing their research. Particular mention was made of the ubiquitous VR Juggler, perhaps the most popular CAVE software solution of its time, and of its growing incompatibilities due to age and stagnated development. Much like RUIS, the Immersive Render System IRS presented here is an independent response to this same problem, at the same time.

2.2 Human factors and fatigue

StarCAVE (DeFanti et al., 2009) is notable in that it tackles the issue of limited vertical field of view inherent in many systems. It implements 18 rendering PCs with fully 34 projectors for 15 screens arranged in a 3 high pentagon arrangement. Despite limited floor space for movement the work demonstrates the significant benefits of providing users with a large space for interaction, thus negating the highly undesirable characteristics of a head mounted display. Issues such as claustrophobia and motion sickness due to an unstable horizon, latency, close-focussing and other limitations are well-known amongst users of HMDs yet commentary regarding how this might impact studies is often notably absent.

Analysis and study into the effects of accommodation and vergence in IVE's demonstrates that most studies focus upon positive parallax where objects are rendered as if they are behind the display (Bruder et al., 2013). Negative parallax is shown to have much less impact on accuracy and large IVEs with plenty of space have a high accuracy compared to other systems. This is a key justification for a preference of CAVE and Octave environments for research and professional environments. A subsequent study expands upon these findings and recommends that spaces of 6-7m around a user should be implemented to ensure optimal perception of rendered objects (Bruder et al., 2015).

Bachynskyi tackles immersive interfaces from outside the discipline, by looking at muscle co-activation clustering (Bachynskyi et al., 2015). In this extensive work it is proposed that understanding of what is happening in the physical body will permit the design of more effective and less fatiguing mid-air interfaces. There is a strong focus here upon data-driven research and of the accumulation of a large dataset that can be reused to inform future works. It is this type of data-rich, quantitative work that offers the best platform from which to continue research consistently.

Wang develops a world building application using a head mounted display and Unity (Wang et al., 2013). The qualitative results from the small sample group suggested an overall positive experience with the system. Wang also observed that no one made mention of fatigue during the sessions of up to 90 minutes. Although Wang offers no explanation for this it is notable that his system uses World In Miniature navigation methodologies which reduce the need for large user movements and there may be a correlation worth investigating. Although focussed upon touch and finger interfaces this study demonstrates that effective bimanual design may provide another means by which to reduce operator fatigue, at least by comparison with a unimanual interface (Jiao et al., 2010).

A healthy human does not exhibit rapid fatigue from common activities such as walking or drawing or typing or reading. Why then should such fatigue occur so quickly when utilising an Immersive Virtual Environment (IVE)? The fatigue assumed inherent in an IVE appears to be neither consistent nor assured and occurs sporadically across all the papers reviewed. A number of authors presuppose that users will find motion fatiguing and yet a number go on, often in surprise, to relate that fatigue was not experienced. Perhaps it is only specific scenarios that cause such fatigue. Although not investigated in this study it is evident that immersive fatigue is neither assured nor a mandatory corollary of working with an IVE and it seems likely that future research may yet identify and eliminate it.

2.3 Theories of user interaction

Bertrand examines how well dimensional symmetry within an IVE transfers to real world situations when the IVE is used for skills training using either a 3DOF or a 6DOF system (Bertrand et al., 2015). It was noted that although 6DOF interaction training produced higher levels of ‘precision’ and thus effectively more accurate results from training, that by contrast it also resulted in users reporting that they ‘felt’ notably less efficient and effective than for the 3DOF system. This is a clear demonstration that user satisfaction and performance are not always correlated and shows a clear and urgent need for more quantitative research alongside qualitative in order to ensure sound findings.

Generating a desire to utilise novel, new bimanual interfaces can be difficult and even when it is possible to clearly prove the advantages of such systems a significant number of users still retain their preferences for traditional input methods (Yang et al., 2012). Investigating a hybrid technology incorporating both monoscopic and stereoscopic input elements and comparing it to monoscopic only and stereoscopic only systems showed significant advantage when including stereoscopic interaction (Bogdan et al., 2014). The intention was to prove the hybrid system superior to both but it was shown that the stereoscopic system had inherent advantages in some areas and could even perhaps be advantageous in all areas if jitter and selection methodologies could be improved.

Studying display fidelity it is apparent that researchers need to consider what they are trying to achieve because varying system topology is not directly comparable and can have a significant impact upon results (Bowman et al., 2012). Proprioception cues may also be integral to the successful implementation of a bimanual system (Veit et al., 2008), or it may be that bimanual interaction compensates for when such cues are missing (Capobianco et al., 2009). This contradiction in findings may be application related or it may simply be that in order to leverage the maximum potential from users they may require training.

The popularisation of 3D User Interfaces is a relatively recent development. Perhaps the most well-known of these devices is currently the Nintendo Wiimote. With such devices there is a tendency to implement them in a naturalistic way without considering alternatives, or to assume that a 1:1 mapping of motion to action will provide the optimal interface. Whilst in some cases this is indeed optimal in others it is far from so and different implementations should be considered (Bowman et al., 2012). It is often technological consideration rather than functional needs that drive interface development and work is still needed in order to bring these elements together in harmony (van Rhijn and Mulder, 2006). It is also uncertain what the impact of requiring users to operate using more than 6DOF simultaneously is and studies have yet to address this in terms of efficiency rather than effectiveness. Nonetheless it can be argued that spatial navigation at least is inherently a 7DOF problem and should be addressed as such (Stannus et al., 2014). With basic process still debatable most studies focus upon quantifying motion into time and precision, however there are other ways in which to measure this data. It is also possible and perhaps even more useful to quantify co-ordination in terms of movement efficiency and doing so demonstrates that a free-moving device, whilst faster results in a more uncoordinated trajectory (Zhai and Milgram, 1998). Clearly it is important to be cautious of assumption when considering naturalism in design of 3D User Interfaces.

2.4 Interface solutions

The bat (Fig. 1) is perhaps the very first free-air controller, certainly the first that is a work of scientific research (Ware and Jessome, 1988). It is seminal in introducing the concept that whilst a light pen with 1:1 mapping results in fatigue the bat, with mouse-like relative mapping, was not observed to cause significant fatigue over the course of use by a 'large number' of casual users and thus 1:1 mapping should be implemented with caution. This same work also importantly concludes that the most useful additional tool is a z-axis flip and that the low-resolution of tracking is easily mitigated by a 'gear shift or gain controller'. These findings demonstrate that specific tools have a more significant impact over the long term than generalised technologies. Many studies have ignored the findings of key research such as this. A lack of standardised immersive interfaces, leading to a need to repurpose commodity gaming devices (Scerbo and Bowman, 2012) seems to impede attainment of research goals.

Fig. 1 - The 'BAT' inertial free-space controller

Research density into the feasibility of immersive CAD continues to increase in recent years and genuinely useful data regarding interface implementation, such as the issue of tools breaking the sense of immersion and interrupting the flow of design, are beginning to become known (Wang et al., 2013). These are examples of problems that do not exist in a traditional environment but are critical to wide scale adoption of an immersive one.

The DesignStation system supports mixed 2D and 3D, unimanual and bimanual interactions in order to perform Computer Aided Design (De Araùjo et al., 2012). The system demonstrated the value of fusing interaction data to improve overall performance.

Pieglass is a circular menu divided into segments and intended for gestural selection using two hands (Rioux et al., 2004). Two-thirds of the initial test group of this study made no attempt to move the menu system until instructed to do so. When modified by the addition of a handle in order to resemble a magnifying glass all of the subsequent test group users moved the menu without prompting. This work supports the theory that most users operate under the assumption that paradigms in a virtual environment will parallel those of the real world and accordingly tend to seek out and react to familiar cues.

In his keynote address Bolas observes that the very definition of a user interface is changing and that our behaviour patterns with regard to an interface are now as much a product of the technology as the technology is a product of our needs (Bolas, 2014). The address goes further to suggest that implementations of technology itself are that it determines what future users expect and seek in terms of functionality and that such a process of invention is in-fact symbiotic.

O’Brien offers up an alternative approach and an interesting anodyne to demands for simplification (O’Brien et al., 2008). Rather than attempting to reduce the number of options the user has in order to improve usability through reducing cognitive load O’Brien theorises that by overloading basic functionality the benefits of an increased tool set will outweigh the disadvantages of complexity.

Li demonstrates that immersive technologies benefit from different types of interface implementation (Li et al., 2015). The research regards 7DOF navigation, (translation and rotation plus scale) and shows that using a linearly multiplied offset offers advantages over a fixed offset, no offset and quadratic (go-go) offsets. Contrary to expectations, parallax effects from targets being drawn in front or behind the physical screen surface have no measurable impact and only the usability of the interface system impacts the results. In this case usability means a linear offset to the hand controllers so that they operate more-or-less 1-to-1 when close to the body but with a significant multiplier when distant. A simultaneous study examines the devices themselves used within the study and concludes that each offers very similar performance (Feng et al., 2015).

Considering the objective of the user rather than designing systems which mirror existing technology is central to leveraging the advantages of virtual interfaces (O'Hara et al., 2014). The use of the ubiquitous mouse pointer within an immersive environment is one such example of a paradigm that transfers poorly and due to the depth element of an immersive system alternatives should be sought (Schemali and Eisemann, 2014).

The GlobeFish and GlobeMouse are novel hardware interfaces used to explore docking techniques in an immersive environment (Froehlich et al., 2006). Although focussed on hardware development the work revealed that separating 6DOF motion into a 3DOF rotation and 3DOF translation component and permitting the user to control these separately improves performance in their docking task. Simply limiting the axes does not guarantee an improvement in performance and in a thorough comparison unconstrained 6DOF bimanual interaction utilising haptic devices showed significant performance advantages over 3DOF constrained interaction (Weller and Zachmann, 2012). In order to understand these conflicting outcomes one must consider the task and objective asked of the user and how the constraint, or lack thereof, benefited them in completing the task.

Works that inform specific knowledge approaches, such as methodologies for selection of multiple objects in a 3D environment, are highly effective at conveying information that can be easily reused. In one study the implementation of an ensemble of complementary tools with a 3D spherical brush is recommended as the most effective idiom for generic applicability to a number of scenarios (Stenholt, 2012) .

The barriers to adoption of immersive techniques for design are high and many researchers are openly dismissive about the potential of such systems. However, even such wary proponents of immersive technologies are frequently surprised by the effectiveness of the technologies they explore and often conclude that developing strategies for managing interaction types is a critical factor for adoption that is far more important than enhancing tools (Israel et al., 2013). Whilst simplifying experiments in order to eliminate users being distracted by the lack of irrelevant features is a logical approach for foundational research the development of advanced tools demonstrates significant potential (Kil et al., 2005).

Surveying the state-of-the-art highlights that key factors affecting current systems are capital outlay and finding the means to fuse the technical and artistic components of a system, the latter being critical for success (Thomas, 2012). With such high requirements to begin researching 3DUI’s for immersive technologies, both in technology and skills, it is unsurprising that good research is sparse. It is rare to find research performed without access to premium technology but it is not impossible to do so as evidenced by one work which fuses a number of consumer technologies to explore the usability of a new type of free-space haptic audio controller (Niinimäki and Tahiroglu, 2012).

Novel approaches that circumvent a need to commit to developing underlying technologies before any useful progress can be made, and thus can focus upon outcome, often offer outstanding results. One such work regarding bimanual hand tracking with only an RGB-D sensor devolves a pair of hands into 26 Degrees of Freedom and computationally calculates finger and palm positions even when obscured (Oikonomidis et al., 2012).

Some argue that the immersive environment requires an entirely different approach to a non-immersive for the purpose of learning how to best utilise it. Within an immersive system the presence of a user is inherently an on-going and continuous interaction that requires the system provide on-going response. For this interactive response system to be effective it must be an inherent part of the system and reverting to traditional text-style help files or static prompts both breaks the immersivity and fails to leverage the advantages of an immersive system. One solution to this is to provide a system that enables the computer to interrogate itself and provide the user information regarding the options available to them at any given time (Murray, 2011).

2.5 Performance improvements

Assuming that bimanual division theory of assigning gross accuracy to the non-dominant hand to gain a performance increase will always be true does not take account of the task and in some instances may be ineffective (Ulinski et al., 2007). Meyrueis introduces a method for deforming CAD objects in an immersive workspace using a bimanual control system (Meyrueis et al., 2009). The work focuses upon the deformation process but relies upon bimanual, asymmetric division of labour as the primary hand is tracked for local object interaction and manipulation, and the offset hand performs command and pilot functions for global positioning and, by proxy, manipulation. This interactivity is assumed however and its impact upon the effectiveness of the D3 system is not considered. Again it is evident that much otherwise good research is not examining techniques in isolation but rather working within numerous factors and failing to acknowledge these multiple impacts.

An extensive review into grasping virtual objects explores a number of visual techniques and strategies for improving the implementation of single handed grabs (Prachyabrued and Borst, 2014). The work observes that although the standard for grabbing is to constrain the virtual hand so that it collides with the object in a manner similar to reality this does in-fact result in the worst performance. By comparison, grabs that penetrate the target object offer amongst the highest performance but are considered subjectively ‘bad’ due to their unrealistic visuals. An alternative is suggested whereby a grasp offer both a realistic and unrealistic component, either a transparent secondary hand shadow to represent penetration or else some form of colour indicator upon the target to symbolise grasp pressure. This results in near-comparable performance whilst offering more ‘realistic’ visuals. For the purpose of this work performance as a metric is preferred above other the metrics and this bias must be acknowledged.

Teather studies the impact upon effectiveness of selecting virtual objects when enhanced with texturing to enhance depth perception based cues and colour based highlighting to enhance positive selection (Teather and Stuerzlinger, 2014). The results demonstrate that texturing had negligible impact but that colour based highlighting offered a significant improvement in selection accuracy. Highlighting impacted time taken negatively with a significant increase. Teather’s test system was a tracked fishtank VR system based around a desktop stereo monitor and from the sample images and his own commentary it seems likely that the system wasn't providing a high level of immersivity and that users were resorting to a point-and-hope strategy. This could explain the increase in time taken along with increase in accuracy since this may correlate with the idea of users waving the selection tool approximately in the vicinity of target space until the colour indicator confirms positive alignment.

A simple but thorough study into selectivity within an immersive environment concluded that greatest selection inaccuracy occurred along the user's view axis (Lubos et al., 2014).

Achibet's virtual mitten paradigm investigates a novel idea for a bimanual controller (Achibet et al., 2014). As a corollary to their core work of haptic feedback they note that visual feedback provided inconsistent results which variously aided or hampered users. It is evident that the impact of visual aids common in monoscopic systems when transferred into immersive systems are not well understood. Incorporating paradigms commonplace in 2D environments into an immersive system does not always yield the response expected and can lower performance (Paljic et al., 2002).

A few studies consider including haptic feedback as part of the selection process and one such concluded that visual feedback was dominant and that haptic feedback could provide an enhancement of this (Basdogan et al., 2000). As an interesting aside it also observed, significantly and conclusively, yet without being able to provide an explanation for it that users associated haptic collaboration with male users and lack thereof with female.

An unusual study into feedback systems for immersive technologies compared visual feedback to haptic and electrical muscle stimulation (Pfeiffer and Stuerzlinger, 2015). The results showed a slight advantage to visual feedback. Another similar study augments visual interfaces with electrically stimulated tactile feedback but is still only at the stage of exploring device development and offers no comparative results or evaluation of effectiveness (Bau and Poupyrev, 2012).

2.6 Bimanual interaction

It was Guiard's seminal work that laid the foundations for all current bimanual understanding (Guiard, 1987). Guiard's kinematic chain model demonstrated that the non-dominant or off-hand was most often used to set a coarsely grained reference in both the spatial and temporal frame. The dominant hand then performed within this frame to provide fine-grain adjustment. Although Guiard's work has critical relevance within the field of IVE's it was not written for nor intended to be specific to this and represents a generalisation of bimanual motion. A more specific example from the same era is a carefully designed experiment into bimanual interaction which partitions the task so that each hand takes on a specific role (Buxton and Myers, 1986). The results remain relevant and demonstrate the performance gains to be had in a well-implemented bimanual system.

Much of the understanding of how to implement bimanual interaction theories in practice for immersive environments is found in publications that occur around a decade after the publication of Guiard's work (Hinckley et al., 1998, 1997). Designing interfaces for two hands is not a simple matter of assuming that a bimanual interface will automatically save time and that such interfaces cannot expect to directly replace 2D paradigms with success. These works confirm Guiard's theories and propose that future work explore theories such as how and when users should switch between symmetric and asymmetric bimanual interaction for best results.

An investigation using a responsive workbench, PINCH gloves and a stylus demonstrates the efficacy of both symmetric and asymmetric bimanual interaction by comparison to unimanual interaction (Cutler et al., 1997). The results are limited to observation.

Researching virtual interaction utilising a two-handed interface (THI) by comparison to a keyboard and mouse interface (KMI) offers insights that, whilst inconsistent in a number of areas, show with reasonable clarity that the KMI was clearly preferred by novices (Seagull et al., 2009). The data supports the proposition that this may be due to a steep learning curve where experts with pre-existing skills find the THI robustly superior. The THI is criticised regarding the complexity of learning to operate 12 DOF control.

This research into a 3D printed two-handed interface offers many comparisons and references to other comparisons (Schultheis et al., 2012). Users were noted to spend significant time trying to adapt to viewport rotation methods rather than upon placement which demonstrates how there are so many factors that can easily interfere with data acquisition in such a complex system. Placement speed was the sole metric and results were overwhelmingly in favour of THI. The author felt this favourable response was in large part due to comprehensive training.

Another bimanual investigation made initial presumptions about how users use two hands and limited the users accordingly (Capobianco et al., 2009). The results were, unsurprisingly as previously discussed for such assumptions regarding naturalism, contrary to their hypothesis. It was suggested that a period of learning would mitigate this but whilst this might help achieve the desired results it more likely it would also serve to hide the truth that constraining bimanual interaction to expectations will inevitably have a negative impact since it forces some users to behave in an unnatural manner. Although the primary results of this work should probably be dismissed they did also usefully observe that depth accuracy was impacted differently to lateral accuracy and that proprioception cues overrode visual cues when the latter failed. Neither observation was substantiated, however other works of more significance have also made these observations.

In evaluating egocentric object selection Poupyrev showed that although the Go-Go technique offered advantages over both ray casting and the classical virtual hand paradigm when operating with distant objects that local objects did not benefit from selection enhancement and that a virtual hand provided equally optimal performance (Poupyrev et al., 1998). A comparison of three methods of pointing and selection using a 3D work surface found that using a real hand directly provided fastest performance but highest error (rate and localisation), an offset virtual cursor provided best precision (lowest error rate) with moderate movement time and offset virtual hands were slowest with moderate errors (Bruder et al., 2013). Thus real hands are better for speed but virtual cursors for precision. Why the virtual hand was poorer than the real one despite being scaled to the same size was not known. It is perhaps simple and easy to convey proofs like these that aid new works the most by providing a concept that can be recreated and built upon.

A study into co-operative unimanual working with 60 users grouped into 30 pairs showed that these paired techniques can provide increased performance in difficult scenarios (Pinho et al., 2002). The study observed that transition from single user to co-operative was no issue for any user due to careful system design. The principles of this study can be applied also to a bimanual system. The key term here is ‘careful system design’: The weak link in the development of an effective bimanual control system is most commonly the way in which it is actually utilised, the ‘interaction metaphor’, rather than the technology or implementation itself which commonly shows great promise (Kunert et al., 2007).

Duval’s SkeweR system investigates the concept of collaborative manipulation within an IVE (Duval et al., 2006). Although Duval initially indicates two bimanual users this is quickly simplified into two unimanual users, each controlling one of two ‘crushing points’ that work in unison. Each user has a ‘tracker’ in one hand and a ‘trigger’ in the other thus this work is effectively a single bimanual system. Implementing a true bimanual system is orders of magnitude more complex than a unimanual system and Duval ends by commenting that he would like to investigate three or more ‘crushing points’ and postulates that four or more would be difficult to implement due to the constraints that would occur.

Mine performs two studies which acknowledge the limitations of a non-haptic environment and consider the ways by which these limitations can be offset (Mine et al., 1997). The study into local object selection, docking and manipulation strongly confirms the preference and optimal effectiveness of local interaction when exploiting proprioception for localisation in an immersive virtual environment. The subsequent study discusses the preference for and benefit of ‘hand-held widgets’, virtual tools that can be held and moved as opposed to fixed in-world tools. The evidence is unequivocal regarding the positive impact of these however Mine overlooks the fact that such widgets are not merely mobile but in-fact leveraging the bimanual skills inherent in the user. This is an excellent example of how easily the benefits of bimanual interaction are not only overlooked but classified as something else entirely.

A novel study utilising two mice with scroll-wheels showed improved performance and preference for bimanual scrolling (Yin and Liu, 2010). This once again reinforces the notion that, despite a lack of real-world implementations, bimanual action is both preferable and can offer performance improvements in many aspects of interaction.

An interesting study into bimanual interaction notes the difficulty of translating physical motion into virtual motion where the virtual object has fixed dimensions but this cannot be reflected in the real world where there is in fact no object (Garcia-Robledo et al., 2009). The solution given is to treat each hand separately and then computationally calculate an equivalent location in the virtual world. The same technology is reused in another study regarding haptic feedback and perceived weight which confirms that users are around five times less sensitive to virtual weights than to real ones (Giachritsis et al., 2009).

Another way to treat the concept of bimanual interaction with a single object concept is with a handlebar metaphor (Song et al., 2012). This study is executed largely through qualitative review but does demonstrate that such a concept is generally effective with the caveat that some users required training to achieve competency. The study also introduces the notion of using a cranking motion to perform axis-constrained rotation but the benefits of this seem inconclusive.

Investigating collaborative interfacing it was shown that the biggest negative impact on results came from a lack of haptic feedback so that two users were unable to synchronise their actions accurately (Ruddle, 2002). This issue is similar for bimanual interaction when handling a single object and it remains unanswered whether such computational solutions as the ‘handlebar’ paradigm are sufficient to offset this problem.

A thorough investigation into the impact of increased dimensional symmetry in a virtual environment intended to train users in surgery shows that an increase correlates to a real-world improvement (Bertrand et al., 2015). The work also demonstrated that constraining the axes which the user operated on adversely affected performance whilst causing the users to perceive their performance as improved. Bimanual interaction with full-freedom provided the most accurate and effective results but not the highest satisfaction.

2.7 Computer aided design for immersive systems

Bowman introduces the idea of interactive visualisations (Bowman et al., 2006) in a simplistic form, distinct from consumed, or experienced visualisation. To avoid confusion this work henceforth assumes that traditional Immersive Virtual Environments (IVE) may or may not have interactivity but Interactive Immersive Virtual Environments (I2VE) mandate some arbitrary high level of interactivity with the environment.

Rahimian's highly comprehensive study into the advantages of an immersive design system by preference to traditional methods of architectural design offers a large number of positive affirmations (Rahimian and Ibrahim, 2011). Such a system, utilised over an extended testing period, is shown to reduce design time, increase collaboration, increase cognitive throughput and provide many performance increases. In particular it is shown to increase experimentation diversity within the early design stages resulting in less traditional resultant models and more thorough examination of solutions. The impact of the medium upon the design results is highly evident and quickly becomes clear that traditional pen, paper and modelling mediums by their very nature place constraints upon the range of solutions. It is common to think in terms of the limitations of Virtual Reality design systems but Rahimian's work shows that traditional methods can be more limited.

Rahimian's study is one of few high quality research works yet available for the field of true I2VEs is still young. MakeVR (Jerald et al., 2013) is another often cited work but utilises only a small participant sample of four for evaluation and does so with minimal constraint. It offers little insight beyond a declaration that the system is 'engaging' but does at least help confirm that such systems are of interest. In terms of hardware technology MakeVR is deployed upon a 3D monitor and no mention is made of head-tracking.

CaveCAD (Hughes et al., 2013) was developed for the University of California's StarCAVE. Like MakeVR it utilises a very small participant sample of four and whilst more rigorous than Jerald's study is still subjective. Hughes observes that the participants who were experienced in CAD performed better than those who were not without analysing whether this was impacted by the unique qualities of an immersive environment. Hughes' participants also resorted to commentary desiring an undo feature and to mentioning arm fatigue, both of which suggest that flaws in the experimental process are likely obscuring data of interest.

With twenty years of publications and working for Walt Disney Imagineering Mine's research into immersive CAD is refreshingly candid (Mine, 2003). Mine casually acknowledges that, "Building a real-world immersive 3D modelling application is hard." before succinctly suggesting that traditional methods, whilst acceptable, inherently discard a "wealth of spatial information" and impose constraints that reduce the potential effectiveness of the design process (Mine et al., 2014). Mine acknowledges the issues and limitations of interacting in an immersive workspace, notably fatigue and that leveraging controllers such as the Wiimote and Razer Hydra not designed for such use introduces additional constraints and problems. Mine's work contributes positive reinforcement to the idea that there is no good reason to enforce 1-to-1 mapping between user motion and avatar action within an immersive virtual environment. Mine's work also demonstrates that by focusing upon the goals of minimising user energy expenditure and maximising user comfort rather than defining a particular hardware specification or technology results in an effective outcome. Mine is clear that although developing a system is very difficult there is an absolute necessity for the researcher to divorce themselves from expectations and assumptions if they wish to achieve a significant outcome. Mine demonstrates that carefully avoiding common but unrealistic expectations whilst embracing thorough, scientific, critical analysis is essential for success within the field of immersivity.

An immersive CAD system using limited resources based upon the Open Source Blender software yet still offering stereoscopic view and both head and hand tracking is developed (Takala et al., 2013). An attempt is made at evaluating the system but the qualitative results offer little insight beyond the system being 'fun' and 'intuitive'. Commentary suggested that the head tracking was ineffective although whether this was due to the limited viewport (a single monitor), limitations in the tracking technology (Sony Move) or some other factor is not investigated. Eye fatigue is also mentioned however which does indicate a possible tracking issue. The paper asserts that much of the technology for Immersive Design already exists in many homes and none of it need be prohibitively expensive.

An economic study into immersive CAD using a Leap motion and Oculus HMD exhibits a significant volume of positive user feedback but little critical analysis (Beattie et al., 2015). A similar scenario is true for another study developing an immersive technology for architectural design which makes for interesting reading but fails to address impact (Anderson et al., 2003). It seems common to find systems like these which are not formally evaluated at all but instead gush over the numerous advantages to designing whilst immersed in a 3D environment. Such exultation does nothing to address the question of how the precision and tools needed to make such concepts professionally attractive can be successfully implemented (Meaney, 2000). ScultUp is also software designed to permit creation in an immersive environment (Ponto et al., 2013) and although the results appear visually impressive there is little but anecdotal evidence regarding the satisfaction of the four users that it was tested upon. There is also no mention of the underlying technology at all and thus, even though reception was positive, replication of this work seems unlikely.

Despite a number of poor quality or impossible to replicate studies that exist within this field there are some few convincing works with genuinely useful outcomes, and much of this work is recent. One study into bimanual interfacing for mechanical CAD concluded that specific forms of input technology are not necessarily advantageous and demonstrated how adding extra functionality to an existing device was actually far more useful in developing an effective system (Fiorentino et al., 2010). Specifically, using a trackball, a non-linear rotational acceleration system was created and the ability to perform an object ‘flip’ added. Although not the objective this is an example proof that it is better understanding and use of existing interfaces and methods of interaction that is the most often needed conclusion data rather than examples of novel new hardware or software.

Another comprehensive study into the feasibility of immersive, in-world model-editing of CAD objects sought to identify and improve the weaknesses preventing adoption (Bourdot et al., 2010). They found that including and refining haptic feedback offset much of the inaccuracy problems but failed to address the needs of different CAD users.

Developers of immersive systems seem to conflate all CAD users into a single amalgam, and yet there are so many different forms of CAD. Both an architect and a mechanical engineer might use CAD systems but their needs are so very different that entirely different software is popular for each. Perhaps this returns full-circle to a key issue stated near the start of this literature review whereby Bowman identified a need to specialise research by making it application driven rather than generalised. By trying to perform generalised CAD VR research it is possible that the outcomes simply have no specific value to any specific user. And by extension, this lack of specificity then denies researchers the opportunity for iterative improvements through real user feedback, because there are none.

2.8 Medical applications

Unlike the field of interactive CAD, which is still largely emergent, the medical industry has been utilising virtual reality equipment and interfaces in practice for some time. The medical field is, perhaps, the anodyne to generalised research in that it has clear requirements and definitive applications. However, although practitioners routinely use bimanual and virtual reality equipment there is still a great deal of confusion about how to best use the technology. In one study video playback is compared to stereoscopic, immersive training and, perhaps unsurprisingly to a practitioner of virtual reality rather than medicine, the results are that neither system offers a clear benefit over the other in terms of quality of training (Harrison et al., 2017). This is perhaps, somewhat sadly, an excellent example of interdisciplinary communication failure. The true advantage of virtual reality lies not in mere stereoscopic reproduction, which has been available for a very long time indeed, but rather in the interactivity of allowing the operator to explore the potential virtual space and thus build multi-dimensional insights.

Ideally, medical training therefore necessitates simulation that is as realistic as possible, with as transparent an operating space as possible whilst offering the opportunity to repeatedly practice a technique until perfection. Virtual reality tools can aid by offering practice simulations that can also evaluate user performance. Of particular relevance to this work it was recently observed that the motion of the non-dominant hand had a significant effect on dominant hand performance (Zahedi et al., 2017).

It might be assumed that most works would consider bimanual interaction, but in fact a surprising number simply ignore the possible impact of this or simply assume that bimanual skills will transfer when bimanual controls are offered. Indeed a typical example of this is one study where a Leap Motion hand-tracker, a device that inherently supports two-hand tracking is nonetheless used uni-manually and compared to a traditional stylus interface (Wright et al., 2017). Considering the previously mentioned study which shows that the non-dominant hand affects the results of the dominant, a trend other studies have shown before, is it appropriate to compare devices unimanually and expect results to project proportionally to a bimanual scenario?

Although it is tempting to think only of physical, haptic controllers when considering medical applications, since these seem naturally enough to be reflective of a surgical procedure, not all medical operations necessarily benefit from such an interface. Indeed a bronchoscope, when used ideally, would offer the operator no haptic feedback since it should not be colliding with internal structures. Existing technologies in use have a very low success rate for this and alternative systems that provide hands-free control have demonstrated a potential for improvement (Khare et al., 2015). It seems likely that this is an area in which virtual reality has much to offer since great advances are being made in free-space control systems.

2.9 Experimental frameworks and the jerk metric

Surgery requires a better metric of analysis than simple target-oriented and time-oriented tasks. For that reason the jerk metric has been adopted as a powerful tool for evaluation of medical performance studies and has been shown to correlate well with psychomotor skills evidenced in virtual reality (Mohamadipanah et al., 2016).

A number of experimental approaches to jerk metric computation are rigorously explored and evaluated in order to find a quantifiable evaluation method to replace traditional qualitative observation (Estrada et al., 2014). Motion based metrics were found to be strongly correlated to existing and accepted traditional grading assessments of skill in surgeons performing minimally invasive surgery.

A well-implemented jerk metric should be a measure of signal shape rather than duration or amplitude, it should be sensitive to movement changes, consistent, robust and be useful in practical application (Balasubramanian et al., 2012). The jerk metric itself can give different results depending upon how it is applied. In order to address this research proposes a dimensionless, squared form of jerk which is independent of amplitude and duration (Hogan and Sternad, 2009). Different ways of measuring jerk may be applicable to different scenarios but the most common methods, which rely upon derivatives, are sensitive to noise and are sometimes rejected as performing inaccurately. In a comparison of techniques spectral arc length measurement is shown to offer improved performance by comparison to log and non-log dimensionless jerk metrics (Balasubramanian et al., 2015).

In addition to a viable metric careful selection of the experimental framework itself can offer improvements to the robustness of the evaluation. Evaluation within Virtual Environment is often limited to usability studies or task-based studies but a few works have attempted to define a framework. Such frameworks may focus upon evaluating the environment itself (Lampton et al., 1994) or may extend this to include suggestions regarding good methodologies for evaluating motion (Poupyrev et al., 1997) but are aware and make it clear that such works are limited in scope. Use of traditional statistical rigour and techniques such as constructing experiments in terms of dependent and independent variables can add significantly to the evaluation process and reveal information not evident in simpler evaluations (Bowman et al., 1999).

2.10 Applications of free-space bimanual interaction

Free-space tracked controllers lacking haptic feedback, such as used throughout this work, are less relevant to surgical research. However, medical procedures such as bronchoscopy minimise haptic interaction and thus there are applications within these fields. Pursuing such medical research is necessarily difficult and limited due to the human factor and thus best left until such concepts are reasonably established and proven under more easily executed trials. A number of environments have requirements comparable to bronchoscopy. Operations where collision poses risk, thus haptic feedback is of limited value, and smoothness and economy of motion are desired, thus justifying the use of a jerk-based analysis could be considered equivalent. This includes operations in space, nuclear disassembly, teleoperation during disaster management and explosive ordnance disposal.

On orbit servicing of satellites is a sector of growing interest, both scientifically and financially, to which a number of control solutions have been proposed and explored that support the merit of free-space control options for specific cases (Flores-Abad et al., 2014). Even despite the limitations imposed by control latency researchers continue to pursue teleoperation options and devise strategies to minimise the issues this causes (Stoll et al., 2012, 2009). And little wonder when the such teleoperated systems so clearly offer improved performance, improved chances of mission success and reduced costs (Fong et al., 2012). The Robonaut 2 is one such tool, deployed on the International Space Station, which is highly humanoid, bipedal, bimanual and utilises elastic systems to offset the imprecision of the human operators force control and reaction times (Diftler et al., 2011).

Explosive Ordnance Disposal (EOD) is a field well-suited to bimanual control and virtual reality stereoscopy (Kron et al., 2004). Key factors of importance are identified to include dextrous bi-manual control (Handelman et al., 2010), and but also numerous varied methods for provisioning haptic feedback to the operator (Graham et al., 2011). The use of teleoperated platforms to move ordnance away from a high-risk area to a low-risk area for controlled detonation is firmly established one of the prime duties of a field specialist but the equipment for this is largely non-interoperable and the lack of standards causes some difficulties (Hinton et al., 2011). This lack of interoperability and notable lack of standardisation is typical of systems that are rapidly evolving and offer potentially exciting bi-manual features not yet available in the field (Tunstel Jr. et al., 2013), such as the Johnny-T and RoboSally platforms(Fig. 2).

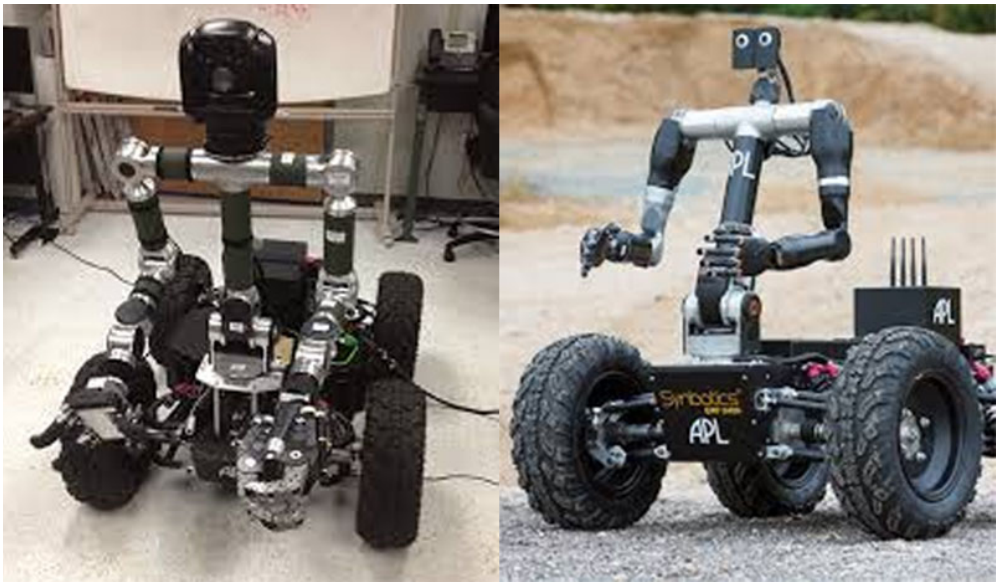

Fig. 2 - Johnny-T (Left) and RoboSally (Right) experimental platforms

Controlling such platforms remains complex however and traditional controllers simply cannot translate human motion, or intent, into the corresponding action successfully without finding ways to better exploit the redundancy of such systems (Wolfe et al., 2016).

The assumption that 1:1 mapping of input or feedback is automatically beneficial or desirable (Mine et al., 2014) (Bowman et al., 2012) could quite logically be extended to this application. It is extraordinarily difficult to map human motion to that of a robot with many degrees of freedom (DOF) (Leitner et al., 2014) especially when they do not align directly to human motion. As technology becomes more complex increasing DOF are highly likely; the Robonaut 2 utilises 42 DOF (Bridgwater et al., 2012). This problem is sometimes termed the Sparse Control Problem and intermediate systems, informed by machine learning and other AI support is considered to be one possible solution (Jenkins, 2008), but this is not without limitations, particularly the potential loss of desirable human characteristics. Equally, when the teleoperated device evidences reduced or constrained DOF by comparison to a human then it becomes necessary to develop strategies to limit human motion, but in so doing there is again a risk of eliminating desirable characteristics motion (Pollard et al., 2002). When the objective is to perform some complex or dangerous task that is difficult for a human and unpredictable in nature then it is human insight and finesse that is primarily required and not so much the arguably limited human range of motion. The issue is compounded further by scale, where it is necessary to interpret inordinately large or small forces in some way and present them to the operator effectively in some abstract manner (Murakami et al., 2000). However, many varied ways of interacting with a virtual environment, each with unique advantages, have been developed and thoroughly proven (Bacim et al., 2013). This knowledge could be leveraged when using virtual control systems intended to direct real systems to gain similar advantages.

For all this wealth of research very little is tested in a bi-manual environment and viability remains low: Technology being purchased for active use often remains solidly in the realm of single-arm, haptic-feedback control, despite the many scientific proofs regarding the potential benefits of bimanual control. For example, the British Army recently took receipt of four such unimanual robotic devices, the Harris T7(Fig. 3), as the primary ordnance disposal tools in a £55 million order (Ministry Of Defence, 2018). It is reasonable to suggest this as an indicator that were an advantageous and proven bimanual alternative available the purchase order would have been different.

Fig. 3 - Harris T7 EOD platform

It seems bimanual research, specifically into how to map bimanual control input to a desired output, has much room for solutions to be found. And the output need not be bimanual, the benefits of a bimanual input mapped to a unimanual output could be substantial. Perhaps free-space control systems, for the very reason they lack the assumed requirements of kinematic constraints, might have some insight to offer the bimanual control paradigm? Is it possible to map free-space input in such a way as to offer fluid versatility that extends beyond a rigid system? What can be learned from such a system?

Perhaps, with careful selection of specific domain, attention to the understanding that has gone before, use of appropriate frameworks, data-gathering techniques and correct metrics for evaluation, then this new knowledge will be of value to a real world application.

3. Design

3.1 Decision making

As this work begins currently popular Head Mounted Displays such as Oculus Rift and HTC Vive have not yet been launched. The facility available is a highly-specialised fourteen screen, reverse projection CAVE with multi-camera infra-red Vicon tracking. This facility is considered state-of-the art and world-class. There are some limitations however, particularly that development upon the software used by CAVEs has stalled and not been updated in many years and that the idea of using two controllers, bimanually, is still an exceptional one and thus there are very limited options for exploring this usage scenario.

The work begins by selecting Open Scene Graph to be the basis for a new engine to drive the Octave facility. The key factors in selection were the wealth of knowledge available for the system, existing departmental experience in its use and a clear assertion within the OSG wiki that it has the latent capability to be used for tracked, stereoscopic systems.

It took two years to develop a fully-functional framework for the Octave based upon OSG. The foremost difficulty was a lack of documentation regarding quad-buffering and stereoscopic projection computation for a multi-threaded, multi-context environment. Much of the information necessary to implement this technology is considered proprietary by the graphics card manufacturers who consider it to be a key part of differentiating their consumer products from their more expensive professional ones. A series of heuristic techniques were used in order to deduce the necessary implementations for a new framework. This work eventually culminated in the Octave OSG framework which was used for the subsequent experiments.

The Nintendo Wiimote was the only commercially available controller intended for true single handed wireless use at this time. Its symmetrical shape, so that it could be used in either hand, along with an available library for interfacing to PC based systems meant it was it uniquely placed as an ideal tool to execute a bimanual study. Its camera based, inertially augmented positional systems were unsuitable for immersive VR use however. Instead, a pair of Wiimote were modified with the addition of retro-reflective spheres to allow the Vicon system with its proven low-latency and high precision, to provide position and rotation data for each hand. The Wiimote was then used to provide button press data.

3.1.1 Octave

The fourteen screen Octave provides an enclosed octagonal 360 degree immersive environment with floor projection and ample space to move freely (Fig. 4).

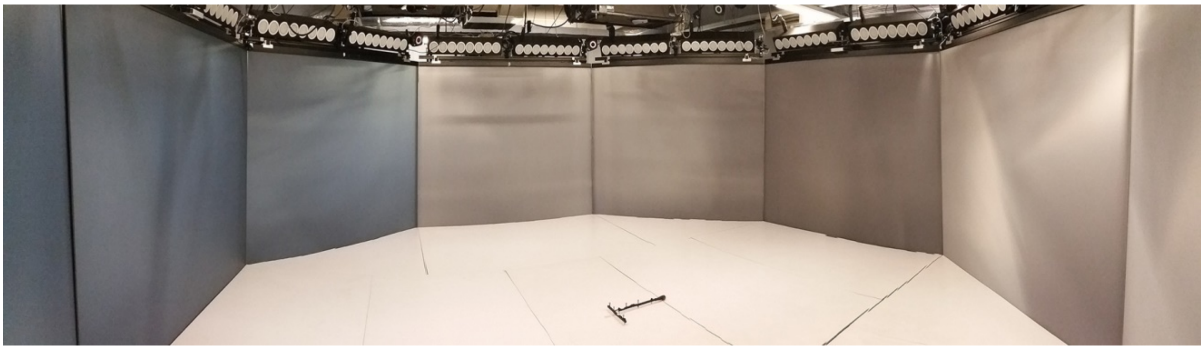

Fig. 4 - Wide-angle view of Octave facility whilst powered down

In order to ensure the best possible evaluation of bimanual interaction in a virtual space the user should not be subject to distractions or cues that break the sense of immersion. A complex CAVE-based system like the Octave offers a number of inherent advantages:

- No limited field of vision. The Octave offers a continuously rendered full 360 field of vision without peripheral distractions or a sense of tunnel vision.
- Reduced distortion. Headsets must use lenses to try to compensate for the close construction of the image projection. The Octave uses only flat screens.
- Proprioception cues are a non-issue because the user can see their own body.
- A distant physical projection point can reduce eye-strain by comparison to a close one.
- No need to calculate horizon data to ensure a steady world since the Octave has a fixed projection reference. This could potentially reduce motion sickness.

The Octave facility is not without weaknesses however. The system is extraordinarily expensive and esoteric and software must frequently be bespoke. It is also not not possible to project an object that occludes the body thus this scenario must be avoided.

The Octave facility also incorporates a Vicon tracking system. This commercial system allows retro-reflective markers to be placed upon objects and tracked with sub-millimetre precision by an array of infra-red cameras. Objects with three or more markers can be composited and then rotation is tracked as well as position.

3.1.2 Immersive Render System

The Immersive Render System (IRS) developed using Open Scene Graph (OSG) to construct and execute the pilot studies is a C++ based collection of source code that integrates with any standard OSG install on either a Windows or Linux based platform. It requires no dependencies and is portable but does require a quad-buffer compliant non-cluster system with VRPN tracker data reporting. Whilst the IRS is a significant work that took two years to produce it is not, in and of itself, of any particular research interest. For this reason only the UML Statechart is provided here in order to avoid bloating this work or transforming it into a technical manual.

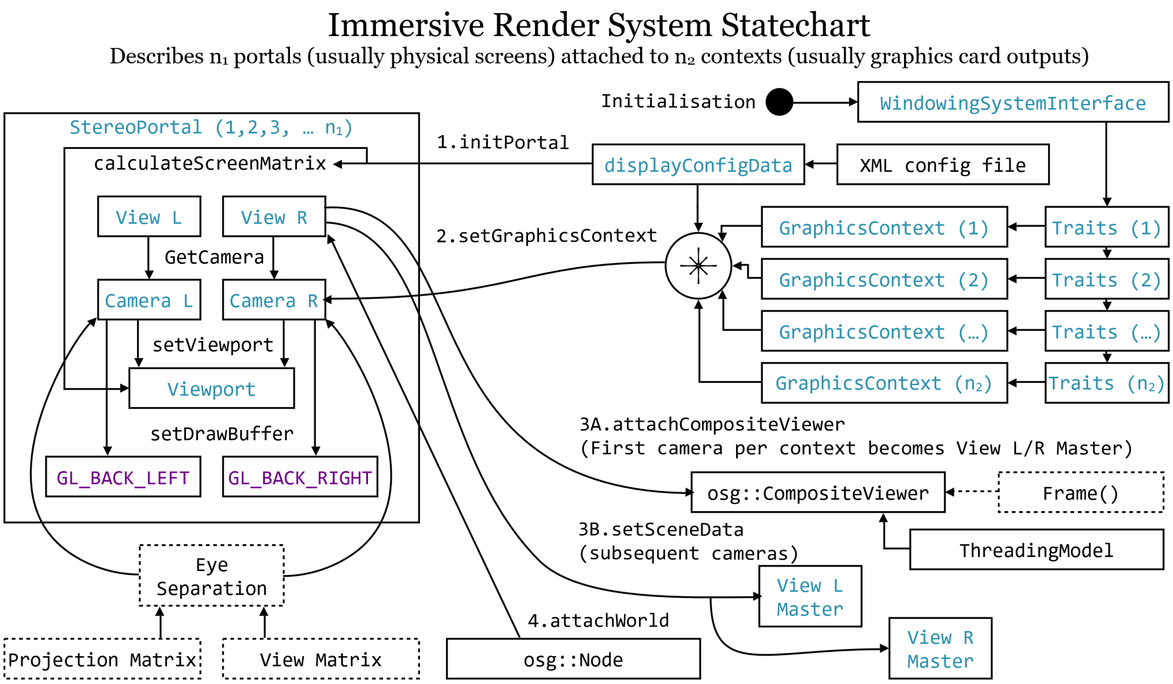

Fig. 5 - IRS Statechart

As can be seen in Fig. 2 the IRS derives from the WSI information to build a collection user definable Trait collections which are used to prepare the Graphic Contexts. The XML user configurable data assimilates this information into the internal displayConfigData format and this is used to build one or more StereoPortals. Each StereoPortal is a self-contained subsystem that manages cameras, projection and buffer allocation. In normal use all StereoPortals are simply assigned to any single standard OSG root node.

The full source code is freely available under an open source license should you wish to recreate or continue the work: See Appendix I.

Prototype demos created using the IRS include a reproduction of the ABB IRB140 robot arm manipulator (Fig. 6) that obeys basic kinematic constraints, and a simple in-situ CAD system that can be used to construct brightly coloured primitive environments, the primary control palette of which is shown in Fig. 7. Links to videos for these are in Appendix I.

Fig. 6 - ABB robotics IRB140 simulation using the IRS

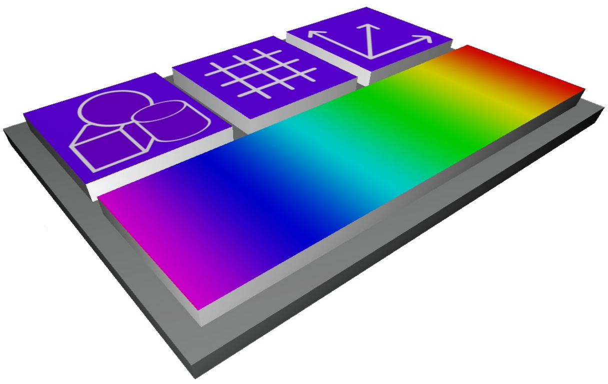

Fig. 7 - Immersive Draw System palette

3.1.3 HTC Vive

Shortly before the final study the Octave facility suffered a hardware failure that rendered it inoperable. It was not possible to obtain an estimate of the time it would take to repair and thus it was decided to rewrite the project to suit the HTC Vive and Unreal Engine 4 in the interim as a safeguard. The Octave repair took a number of months to be completed and included new hardware and a new operating system. Reconfiguring the IRS would have necessitated code changes at this point thus the Vive was used until completion.

By comparison to the Octave the HTC Vive system offers some advantages :

- Laser-tracked sub-millimetre accuracy comparable to a typical Vicon configuration.
- The incorporation of laser-tracked controllers as a unified part of the system (Fig. 8).
- A number of software frameworks that natively support the hardware.

3.1.4 Unreal Engine 4

Whilst a C++ HTC Vive library is available and it would have been possible to rewrite the IRS engine to work with the Vive, examination of information available suggested this would be non-trivial. By chance, the necessity of deciding how to proceed coincided with an opportunity to visit Epic Software who own and use Unreal Engine to develop games for the Vive and Oculus Rift. After an evening in their offices demoing different systems and discussing how this engine could meet the needs of this research, particularly by comparison to Unity, it was concluded that since UE4 is C++ oriented is was a more logical choice than Unity in order to allow direct reuse of existing computational code.

Fig. 8 - HTC Vive controllers (left) and Oculus Rift controllers (right)

3.2 Pilot Study 1

3.2.1 Hypothesis

The principal objective of this pilot study is to test if gathering a significant amount of quantitative hard data will permit this to be explored and analysed in such a way that it provides insight into motion and behaviour. Accordingly, this pilot study takes the form of an explorative process and encourages the user to work with minimal direction or constraint. The primary hypothesis could thus be stated:

> **Ha_0 Given free choice when placing an object a user will evidence a preference for bimanual interaction over unimanual.**

Accordingly, the null hypothesis will take the form:

> **Ha_1 Given free choice when placing an object there will be no pattern of preference for bimanual over unimanual interaction.**

This hypothesis is not intended to be novel but instead is specified purely so that this pilot study can act as both as confirmation study and to validate the design of this study.

3.2.2 Design

The first pilot study proposes a game environment as shown in (Fig. 9). It was observed during development of the IRS that users new to Virtual Reality frequently deviate from the instructions to investigate uncorrelated features. It is expected that a highly structured test in the form of a game will constrain users so that they focus upon the study tasks.

This study uses the well-established ‘pick and place’ paradigm where users are asked to move an object from a start position to a goal position. The objects to be carried are four 30cm cubes of different colours. The target is a magnetic cube which has four of its faces created in matching colours to which the carried cubes will stick when released.

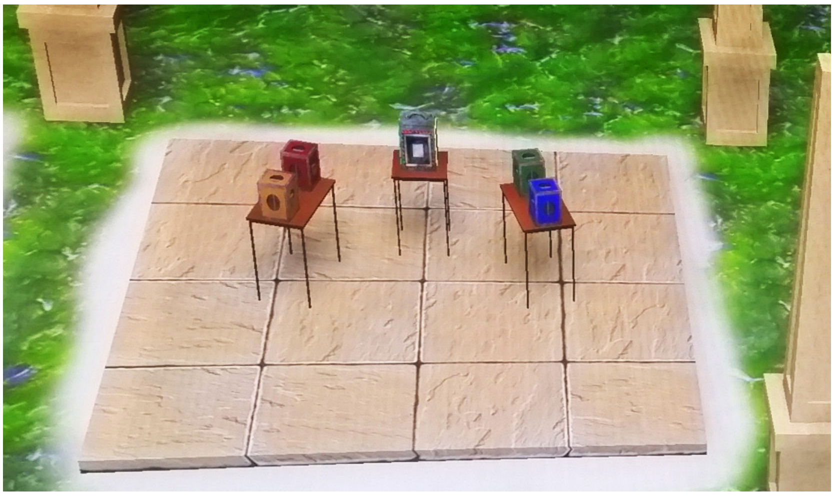

Fig. 9 - Aerial view of the first pilot study game arena

The game takes place within an arena designed to be entirely contained within the Octave so that navigation tools are not needed. This is done in order to ensure that the user remains within the optimal tracking region. Due to the arrangement of the infra-red Vicon cameras, tracking accuracy is biased towards the centre and far regions of the environment. This is by design because both users and applications tend to be biased for forward orientation. The region guaranteed to have best tracking coverage is a roughly circular region of diameter centred 0.4m forward from true centre (Fig. 10).

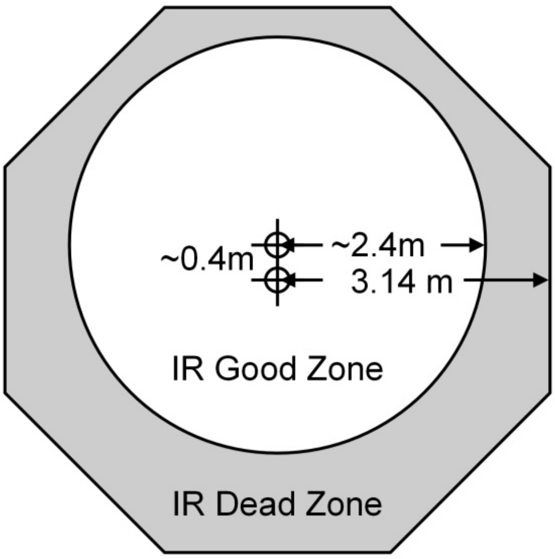

Fig. 10 - Optimal zone for Vicon IR motion detection

The game is designed to fit within this zone whilst maximising coverage for regions that participants are likely to enter and spend time placing or grabbing objects within. The tables are placed to provide visual cues that imply users should pick up cubes from within the space between them. The layout is positioned to ensure there is sufficient clear space within the optimal motion tracking coverage region on all sides of the target cube (Fig. 11).

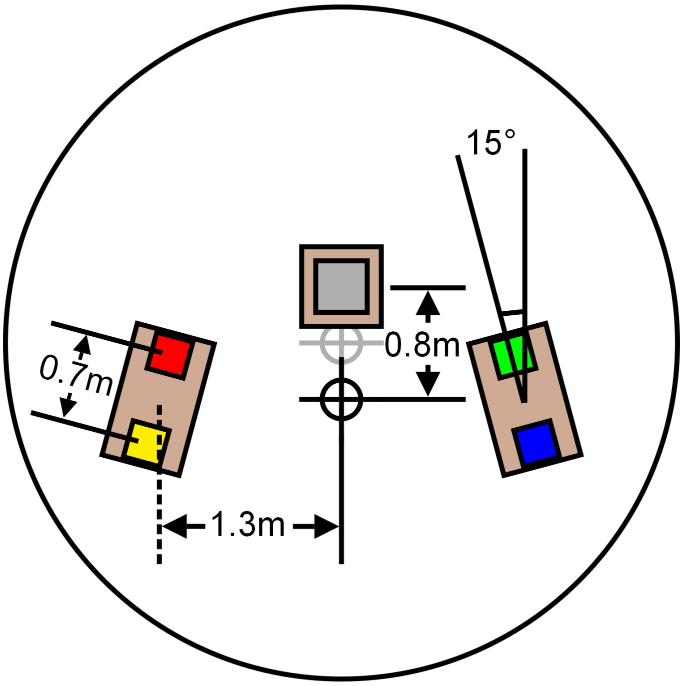

Fig. 11 - Pilot study game layout

The arrangement of the objects is chosen to offer varying pathways in order to promote different types of movement and thus encourage the participants to consider different grab and movement flows throughout the test. It is intended that these cues will stimulate increased participant exploration of the space by comparison with a scenario where the grab and move flows for each colour pair are similar (Fig. 12).

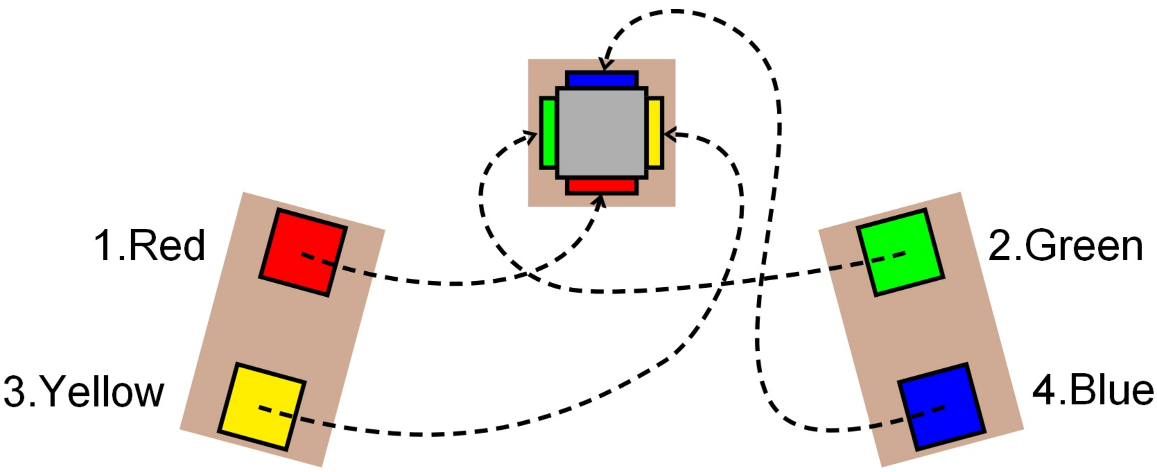

Fig. 12 - Example motion pathways for the pilot study

3.2.3 Implementation

The first pilot study builds upon the IRS by adding four class files:

- `osgSM::scmGame1` provides the game flow.
- `osgSM::scmGameTools` contains reusable game classes.
- `osgSM::scmCollision` is a collection of simulation classes and global functions.
- `osgSM::scmLogger` manages log files and time stamping.

There are around 40 classes within this collection, each of which offers a number of functions. This code is well-documented within the source files and class diagrams can be swiftly generated using Visual Studio tools for the software programmer who requires complete detail. UML State machine diagrams are provided herein however because they are considered effective in describing real time systems and are not easily generated.

Control and input

The user uses a single button to grab objects. The software treats the A button and trigger button interchangeably so that the user may use whichever they prefer (Fig. 13). The system initiates a grab attempt upon the trigger being depressed, maintains this state whilst held and, and will release the object upon the trigger being released.

Fig. 13 - Wiimote highlighting the primary controls

The software operates using a typical game loop, checking for input, calculating updates and then performing a render cycle. The two core processes for gameplay are managing positive user interaction, such as button presses, and managing passive user interaction, such as object collision and motion.

The log system can store data either periodically, in response to user input or in response to world events such as collisions or state changes.

Processing bimanual events

The pilot study software categorises interactions into one of five possible states, termed a GameMode, each of which represents a manner in which the user is utilising two hands :

(0) Select : The user has nothing grabbed in either hand.

(1) PriGrab : The user has grabbed something with their primary hand.

(2) SecGrab : The user has grabbed something with their secondary hand.

(3) DuoGrab : The user has grabbed one object in each hand.

(4) BiGrab : The user has grabbed a single object with both hands.

Primary and secondary hands are arbitrary and are used to track how a user switches from one hand to another. The system operates upon the principle of change-detection, comparing previous to current values, in order to determine how to set and unset object connectivity, GameMode and to update the relevant icon (Fig. 14). The update offset process is used to recalculate origin matrices when the object root is changed.

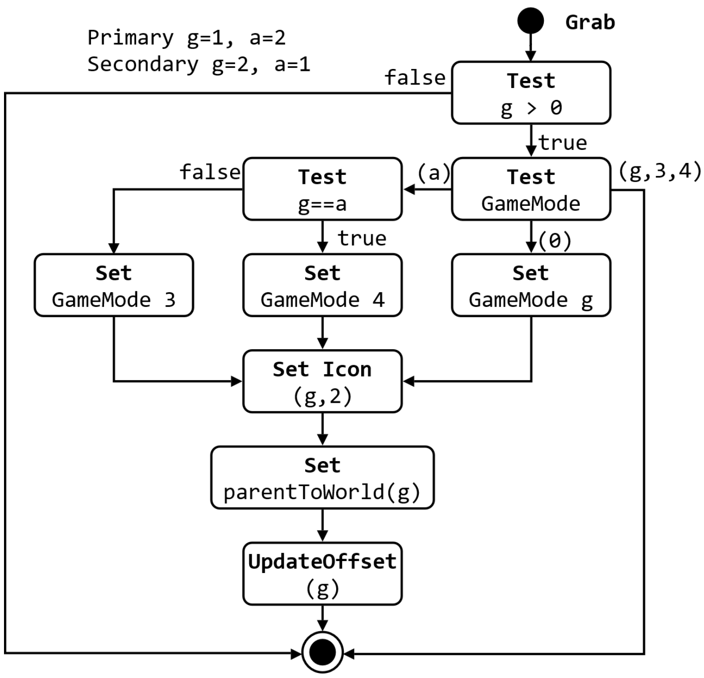

Fig. 14 - UML Statechart for the Grab process

The release process (Fig. 15) follows a similar pattern but also tests for the possibility that the object is in collision with the 'magnetic' target at the moment of release. If this is the case then the object root is connected to the target instead of being restored to the world.

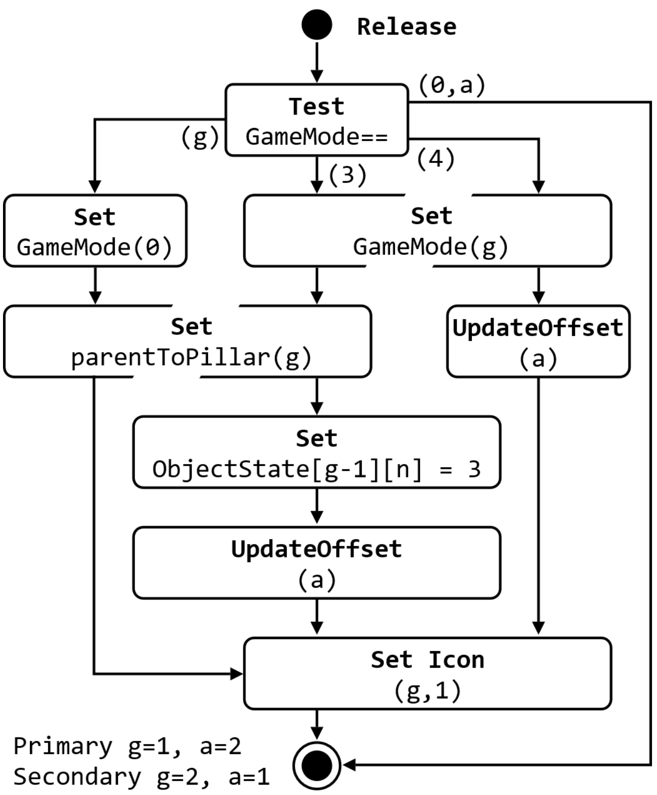

Fig. 15 - UML Statechart for the Release process

Offset updates are computed applying the inverse matrix transform of the hand position with respect to the object position in world space. Parenting also uses matrix re-computation but executes a series of tests to ensure validity first (Fig. 16).

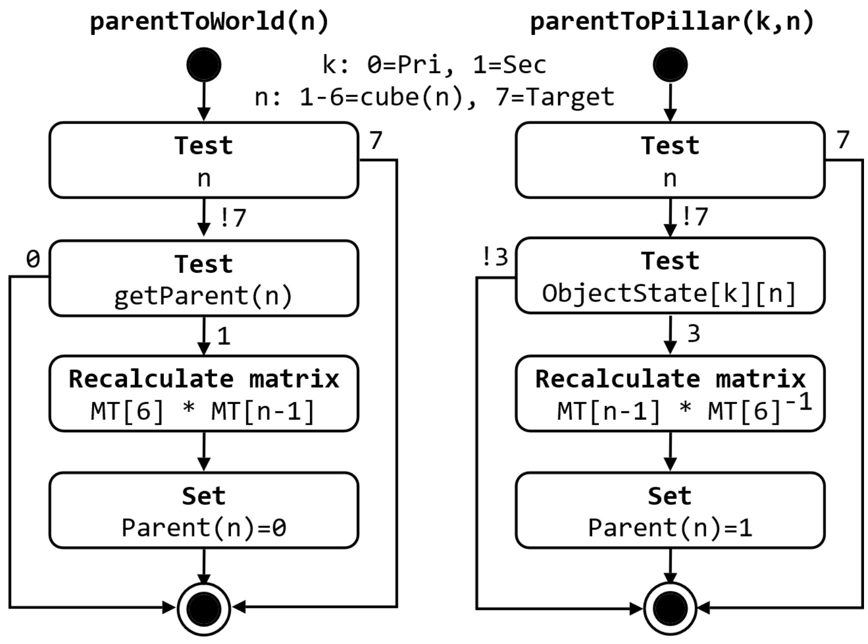

Fig. 16 - UML Statechart for the parenting process

The game system does not use an off-the-shelf physics engine but rather performs tests and discrete collision computation upon specific objects directly. Although this can make development more complex it guarantees accuracy (Fig.17). Typical game and render-oriented physics engines are generally designed with performance and multi-object handling as their prime criteria and do not guarantee precision.

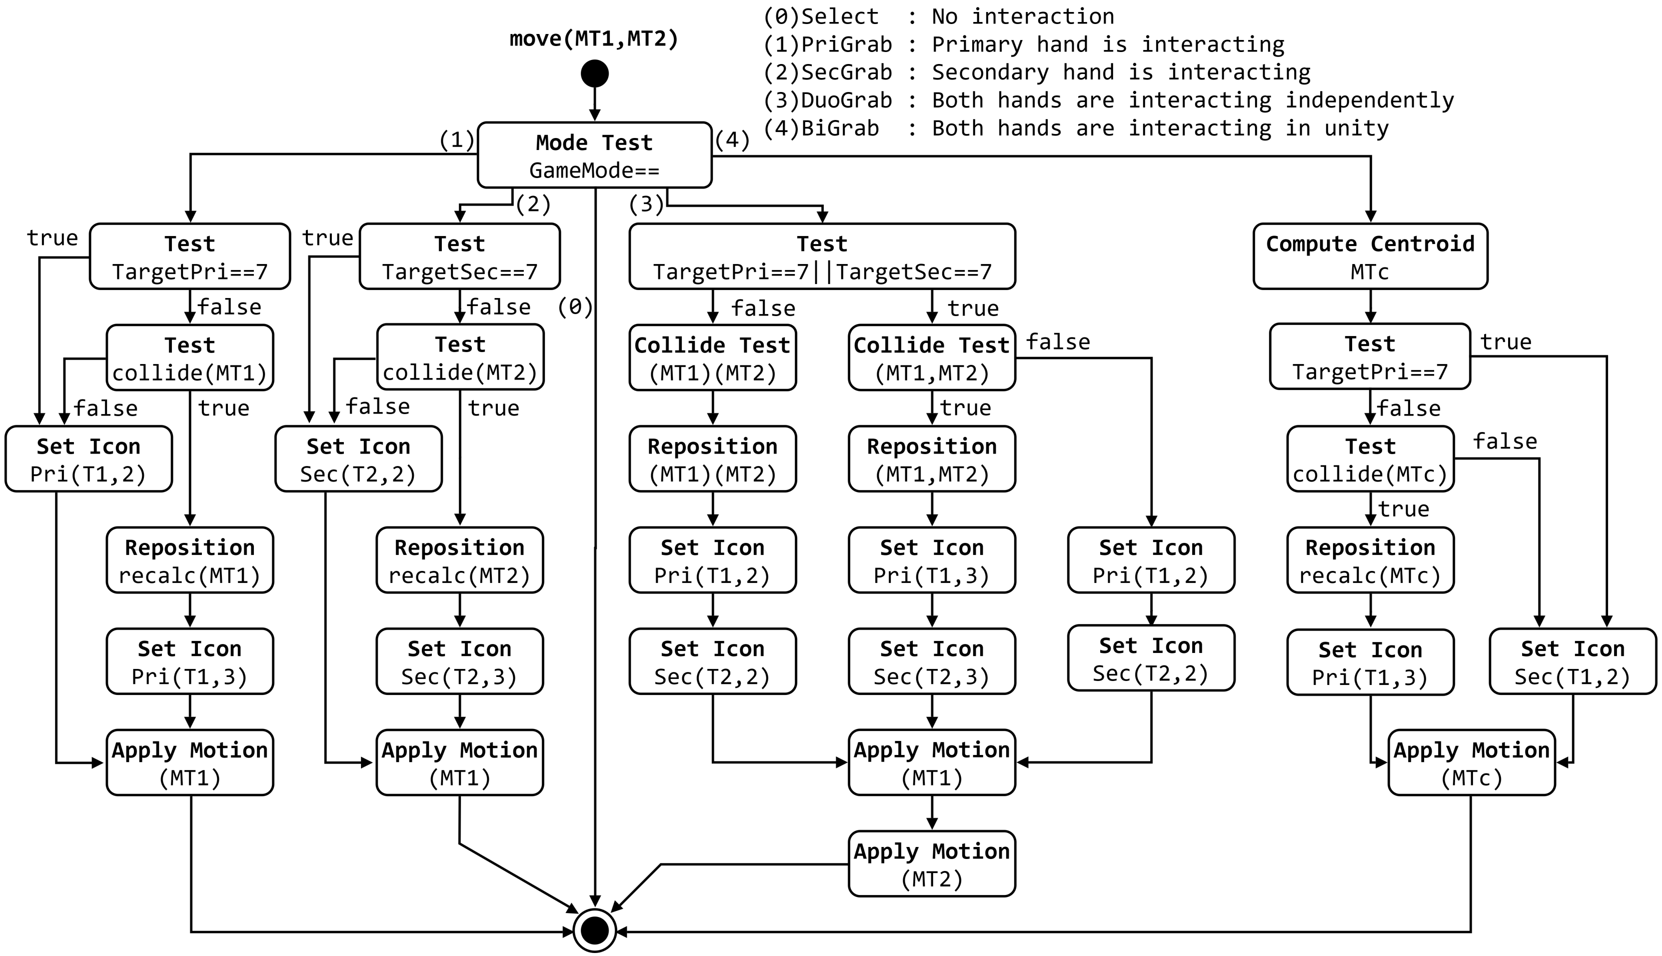

Fig. 17 - UML Statechart for the parenting process

Icons

It is essential to the study that the game display informative icons to the user depending upon whether they are potentially able to grab an object, if they have grabbed an object successfully with one or both hands, if the object is colliding with a target object or for any other state. Fig. 18 shows how the process which tests for change upon every game cycle.

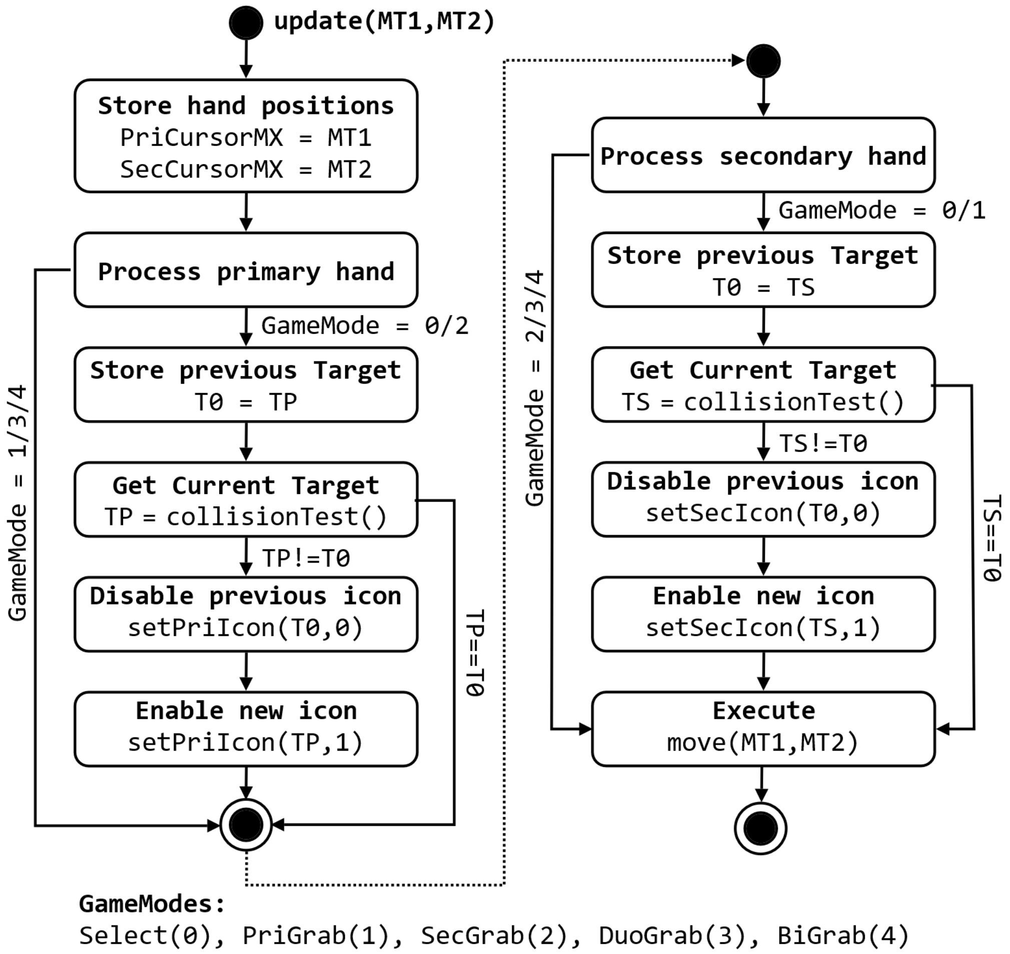

Fig. 18 - UML Statechart for the target update process

Upon calling an icon update an efficiency check is made to ensure that change is required because this is a process that would otherwise be actioned every frame. State changes are maintained even when the icon is disabled for the purposes of data logging and to ensure stability when an icon is re-enabled. Once this check is made update requests are passed to the update process which processes evaluates the permutations of hand state in order to select and activate the appropriate icon (Fig.).

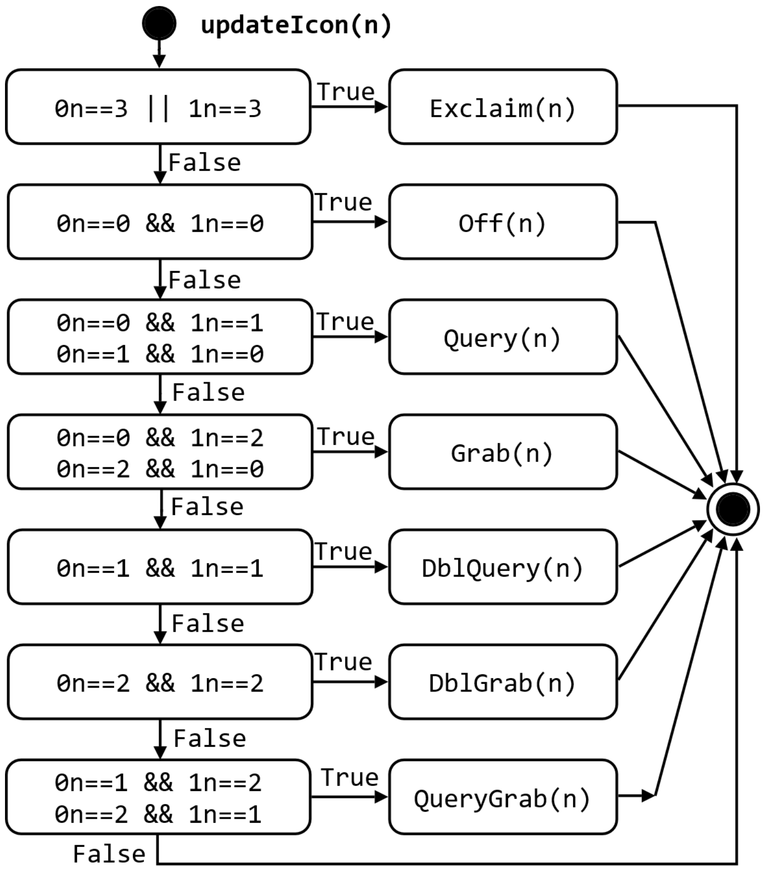

Fig. 19 - UML Statechart for the icon update process

The icon subsystem itself is self-contained. There are more possible icon permutations than game state because the icon subsystem must provide feedback regarding potential conditions, such as the opportunity to grab, and overriding conditions that can modify behaviour, such as collision (Fig.20). It also handles loading of models, manages offsets, rendering, visibility and offers the option for the icons to slowly rotate. Each of the models is only loaded a single time and rendered as many times as necessary in multiple locations.

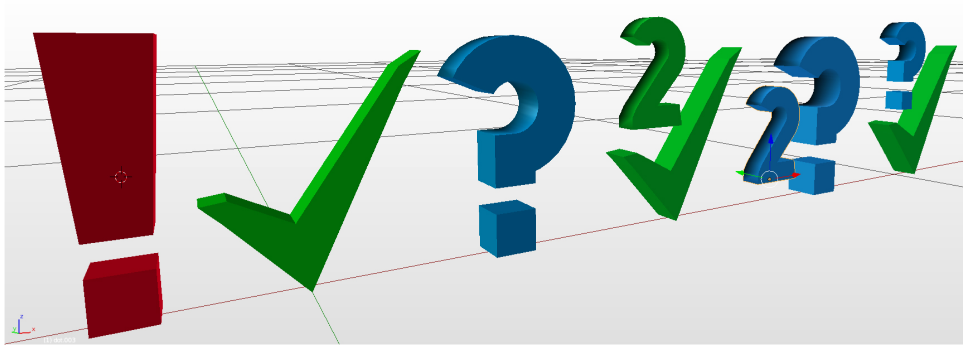

Fig. 20 - Game icons: Collision, Grab, Can grab, Bi-grab, Can bi-grab, Grab + can bi-grab.

3.2.4 Execution

The experiment consists of grabbing, carrying and aligning the four cubes as precisely as possible with the target. The users are encouraged to repeat the experiment as many times as they feel necessary until they feel they have placed the cubes with as much precision as they can. The user is introduced to the environment through informal tutorials then asked to perform the experiment when they feel ready to do so. At the end of each round an approximated score derived from accuracy of placement data is displayed to the user on the in-game billboard along with elapsed time.

Instructions

The following instructions are used by the operator during the process:

1. **Calibration:** The operator verifies the alignment of the virtual world whilst the user is asked to waits outside and read through the instructions.

2. **Training:**

   a) The user is introduced to the system and any questions answered

   b) Unimanual training is introduced where the user is provided with only one controller and asked to pick and place the four cubes quickly and accurately.

   c) Bimanual training is introduced where the user is provided with both controllers and asked to again pick and place the four cubes quickly and accurately.

3. **Study:**

   a) The user is provided with both controllers and asked to execute the experiment in full using any method of control that suits them. They are informed that data is now being recorded.

   b) The user is encouraged to repeat the experiment as many times as they desire until they feel they have satisfied with their performance.

Sample group

The pilot sample group was drawn by word-of-mouth and informal advertisement of the study within the School of Computing at the university. A total of 20 respondents took part in the pilot study. Demographics regarding the respondents were not collected.

3.2.5 Data

Strategy

Similar studies tend to focus upon gathering final placement data, sometimes alongside a time-based metric. This study is explorative so that strategy was extended to collect as much additional data as possible. It was decided to sample all data at every event instance and utilise the findings to refine this strategy of bulk collection in continuing studies.

Table 1 - Data definitions

Data Type	Data Elements
Tracker: Hand 1	7 floating point values (x, y, z, qx, qy, qz, qw)
Tracker: Hand 2	7 floating point values (x, y, z, qx, qy, qz, qw)
Tracker: Head	7 floating point values (x, y, z, qx, qy, qz, qw)
Buttons: Wiimote 1	Only 'grab' and 'release' are used by the Game
Buttons: Wiimote 2	Only 'grab' and 'release' are used by the Game
Game: Object position	5x7=35 floating point values 5x (x, y, z, qx, qy, qz, qw)
Game: Object state	5 objects, each with 7 possible states*
Game: Hand state	2 hands, each reports the current target object
Game: State	5 possible states**
Game: Other	Time stamp, game round, events (start, stop, pause)

**None, Single Hover, Dual Hover, Single Grab, Dual Grab, Hover and Grab, Collide.*

***None, Primary Grab, Secondary Grab, Duo Grab, Dual Grab.*

Each data element contains position and rotation data for every object in the game, position and rotation for each hand and for head tracking data, and entries for all game state and icon data (Table 1). A new data entry is generated every time an action occurs, including collisions, icon changes and button presses. This data would be formatted into a comma separated value format and saved to a plain text file so that it can be imported into many different software options for analysis or transferred into a database.

3.2.6 Results

MS SQL query analysis was used for data evaluation because the author has access to a fast cloud cluster to perform data mining.

The CSV data from the pilot study was sanitised to make it suitable for database processing. This included generating a unique ID per data set, a numerical identifier per user attempt, and adding a unique ID per data row. A series of server queries, links to which can be found in Appendix I, were created to mine the data in search of patterns. The investigation of these queries follows in the next section.

During early mining two bugs in the code of the data logger were found that meant some data must be excluded and this the extent of this first pilot study analysis is limited accordingly. The excluded data is all rotational data, and positional data under the conditions where a user fails to accurately place their object at the target. The software errors were corrected and do not impact the subsequent chapters of this work. The subsequent analysis utilises data unaffected by this bug, such as grab preferences and carry duration, which remain valid regardless of logged location values.

Grab methodology

Queries were generated to mine information regarding user preference of bimanual (two hands, one object) versus unimanual and duomanual (two hands, two objects). The total number of times a grab of each type performed across all games and users was extracted along with the summed time elapsed during each grab process (Table 2).

Table 2 - Query 1 : Grab frequency and duration

Type	Incidents	I%	Duration (s)	D%	D/I (s)
Bimanual	420	23.9	944	39.8	2.25
Unimanual	1242	70.7	1246	52.6	1.00
Duomanual	95	5.4	181	7.6	1.91

This first analysis shows an overall preference for initiating tasks with unimanual control, with almost 71% of grabs being unimanual. Contrastingly, it also shows that the sample group spent a significantly longer amount of time moving or positioning the objects whilst using bimanual interaction.

For final placement a temporal query is built which analyses events leading up to release in order to exclude incidents where the user releases one hand before the other thus recording an unintentional unimanual final placement after a period of bimanual activity. The results of this query reveal that there is a small increase in preference for placing and releasing the object at goal using a bimanual methodology by comparison to the summers interaction preferences for interaction, transportation and placement (Table).

Table 3 - Query 2 : Final placement preferences

Type	Incidents	Incidents as %
Unimanual	93	64.6
Bimanual	47	32.6
Other	4	2.8

Examining the dataset procedurally it shows that users frequently use a succession of unimanual grabs to transport each object then evidence an increased likelihood of switching to bimanual interaction for the final placement stage.

This correlates well with existing works regarding two-handed usage which prove it can provide increased precision in real-world control scenarios. The results show that user preference for bimanual methods increases during the final placement stage. The results do not support an increase in accuracy however, but this could be sue to system limitations regarding the ability for the virtual system to actually leverage the user input fully.

Table 4 - Query 3 : Final placement precision

Type	Deviation	Successful	Missed
Unimanual	9.91 mm	86	11
Bimanual	9.24 mm	43	4

Conclusion to the first pilot study

As a confirmation study it was not proven that users would prefer bimanual control and in fact unimanual control was used more frequently. However, it was shown that users were more likely to switch from unimanual to bimanual during the placement stage than the other way around and to spend more time per bimanual interaction than per unimanual.

This pilot study also serves the purpose of validating the process. In addition to identifying a number of software bugs three critical factors that negatively impacted the outcomes were evidenced which are addressed in the following studies:

Firstly, the data was difficult to analyse because of its size and complexity. Following studies should be more targetted in what they record for efficient and quick analysis. Whilst recording a complete data set does ensure all possibilities can be mined for, the corollary is that mining becomes a sequential process. To be mined effectively temporal processing is often required, in other words identifying overlapping chains of patterns of arbitrary lengths, in order to generate the desired results. By contrast, permitting the software to filter events in real time can completely eliminate this requirement and provide concise and focussed data provide the data sought is well-defined in advance.

Secondly, user interaction was far more varied than expected. Some users would only interact for the briefest possible time whilst others lingered and wanted to repeat the test a number of times. Some users also demonstrated significant uncertainty whilst other confidently pursued explorative goals unrelated to the research. The outcome of this is that it wasn’t possible to evaluate trends with confidence and data gathering progressed slowly.

Finally, the scoreboard at the end of each result appeared to have far more impact than expected. This was not measured and is purely observation but a number of users, upon seeing a score, became competitive and asked what other users scores were in order to best them or at least attempt to. Other users noted that a duration was reported and had to be dissuaded from rushing the experiment. It seems that having a display of information sometimes overrides verbal instruction and thus should be used much more carefully as a valuable tool to support the research outcomes through instruction.

3.3 Pilot Study 2

3.3.1 Hypothesis

The first pilot study showed that most users initially preferred unimanual control but an increased proportion of users swapped to bimanual interaction at the placement stage. This study investigates this further by examining how interaction methodology changes over time. It also extends the study to explore how feedback, both in terms of visual indicators and in terms of simple collision physics, affects placement precision because it was unclear whether users were taking advantage of these facilities in the first pilot study or not.

The primary hypothesis can be stated:

Ha_0 Users will increasingly prefer bimanual placement with repeated attempts

Ha_1 Users will not prefer bimanual placement with repeated attempts

The secondary hypothesis can be stated:

Hb_0 The use of assistive tools will increase precision

Hb_1 The use of assistive tools will not result in an increase in precision

In a similar manner to the first pilot study these hypotheses are intended to support the design and verification of this study and not to offer significant new knowledge.

3.3.2 Design

This second pilot study directly builds upon the first by addressing the weaknesses that were observed. Varying levels of user interaction was a major concern so it was decided to implement much more rigid guidelines and controls for this second pilot study. This includes a structured tutorial and familiarisation stage, a fixed number of rounds of participation, an entirely automated testing process, and a number of smaller alterations to the way in which the experiment is visualised.

Tutorial

In the first pilot study familiarisation was informal. It is possible that this led to some users undertaking the study before they were familiar with how to utilise the tools and thus may not have fully utilised the bimanual interaction possibilities available. It was hoped a well-structured tutorial would mitigate this whilst introducing no disadvantages beyond a slightly longer time requirement per participant.

The tutorial is fully automated. All study participants have exactly the same training and the possibility of the operator bias is eliminated. The tutorial is designed to guide the participant through a series of tasks that introduce all the different ways in which they can interact with the system and offers 5 distinct stages: Primary grab, secondary grab, dual object grab, bimanual grab of a single object and magnetic attachment to target. Each stage disables interaction whilst a script is played back to the user and simultaneously presented as text upon the in-game billboard (Fig. 21).

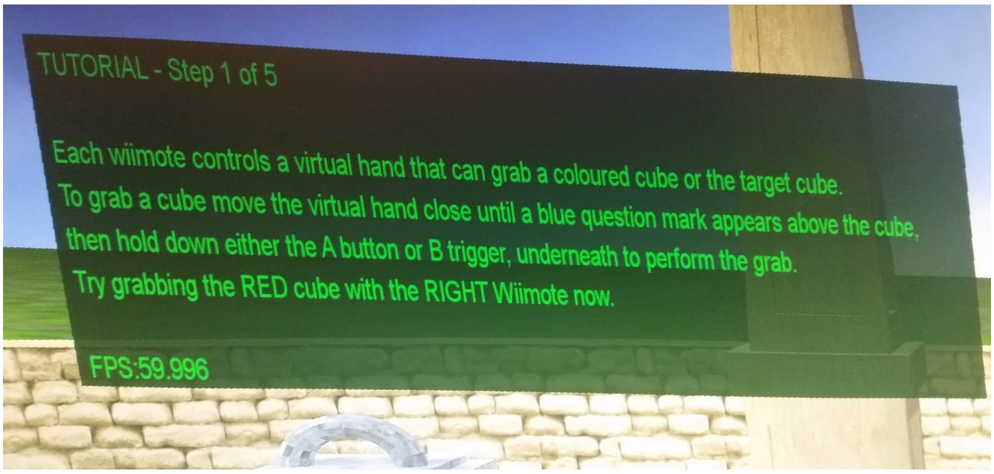

Fig. 21 - Tutorial billboard for enhanced game

Upon completion of each script playback interaction is re-enabled and the system waits for the user to execute the requested task. Upon detection of task completion the system halts interaction whilst explaining the next task to be undertaken. Upon completion of the entire training tutorial the process the automated system allows the user to practice what they have learned and invites them to start the study once they feel ready.

Experiment

The same automation system guides the user through the study itself. The user is asked to pick and place four cubes, as accurately as possible, four times, for a total of sixteen attempts. Each of the four attempts activates a different set of support tools (Table 5).

Table 5 - Game type permutations

Game Type	Collisions	Indicators
A	Off	Off
B	Off	Visible
C	On	Off
D	On	Visible

The order in which the user is offered these permutations is handled by the automation system on a round-robin basis to ensure that the results are free from sequential bias. The automation system also allocates every user a unique ID, every attempt a unique ID and maintains this state across sessions through the use of a secondary XML file in order to further automate the data collection process and ready the output for immediate analysis.

Audio script

An audio script was carefully prepared and recorded utilising appropriate studio equipment, including a large diaphragm condenser, in an acoustically treated environment so that the instructions would be clear and intelligible when replayed. The exact same instructions are displayed upon the billboard whilst being read out. The system integrated the irrKlang C++ audio library to perform playback. The Ogg Vorbis format was used for the audio files because it is free from licensing concerns, unlike MP3.

Environment

A number of alterations were also made to the environment to improve areas that had been criticised by participants or observably could have introduced other biases. Some users had referred to the in-game hand representations as ‘comedy hands’ or had felt they had caused a sensation of vague precision. A simple, natural, high-contrast symbolic representation of a pointing hand as a 3D object proved difficult to come by, most were highly stylised, and so a bespoke solution was modelled based upon traditional cursor colours (Fig. 22).

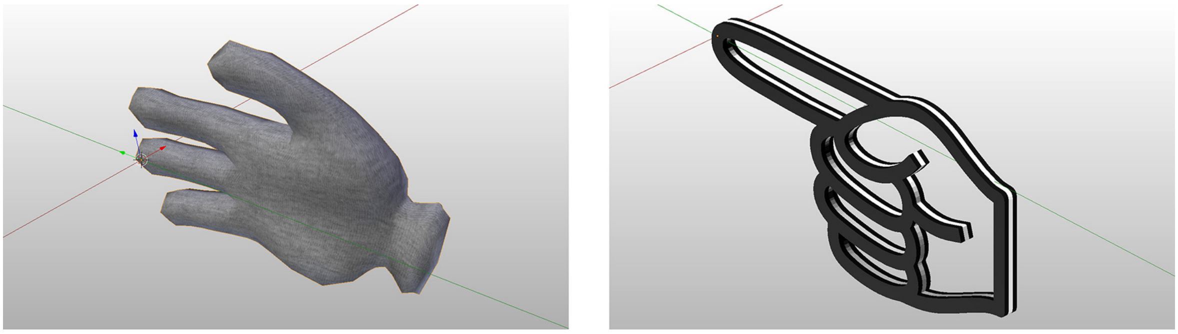

Fig. 22 - Pilot study 1 pointer (left) and pilot study 2 pointer (right)

The world itself for the initial pilot study contained many features useful during the development process but redundant now that the pilot game had been tested. These elements were removed along with the ability to fly around the environment and any other tools that were not directly required for the study. This both improved frame rates and stability whilst eliminating potential distractions from the users environment. The study arena was remodelled as an enclosed area to subtly reinforce this (Fig. 23).

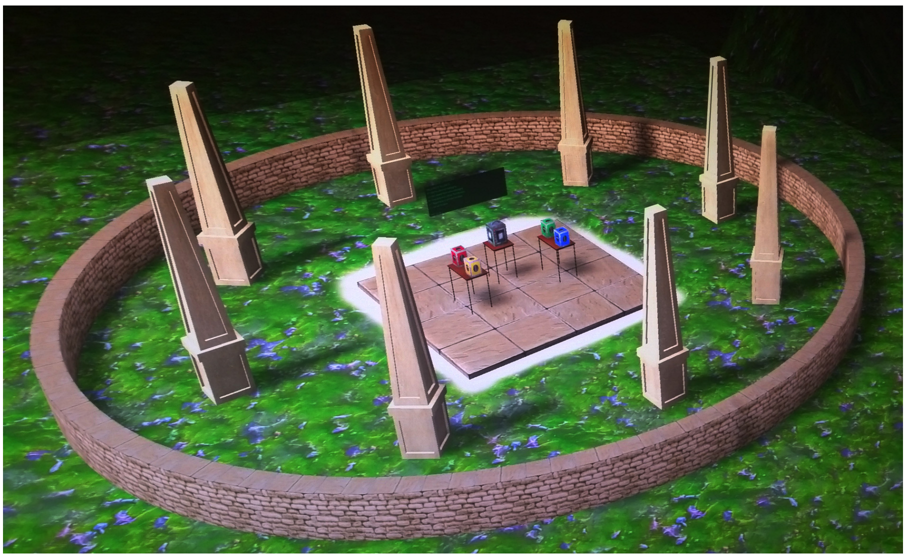

Fig. 23 - Aerial view of pilot study 2 arena

The game arena was also modified (Fig. 24). The two tables and cubes upon them were reoriented from the 15 degree angles of the basic game to being parallel with each other. This design decision was made because it had been observed that the narrowing of the tables seemed to slightly impede users as they tried to move towards the target and the increased complexity of the angled arrangement presented no obvious benefit.

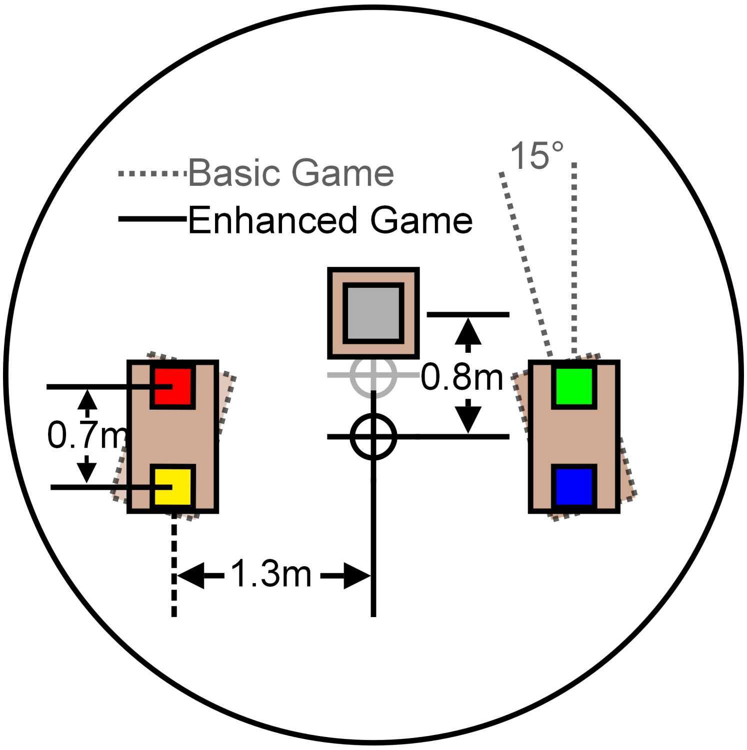

Fig. 24 - Pilot study 2 component layout

3.3.3 Data

The data collection strategy was refined so that only positive user actions, i.e. grab and release, now caused the system to record events. The recording of spurious data from user hands repeatedly passing in and out of objects or coming into and out of collision had previously created a large amount of surplus data.

Data collection was also automated using the same system that managed the tutorial and game execution. Previously data was had to be sanitised and combined manually which was tedious and could easily lead to errors. With the automated system data entries, at every stage, maintain unique identifiers, user identifiers and round identifiers. This was, in part, achieved through the use of a second log file to maintain overall data about the number of entries in each experiment so that unique IDs could be contiguous.

Timestamp accuracy for all log entries was improved to the millisecond level and the log file was made fully asynchronous to eliminate the small risk of write delays.

In a similar fashion to the first study a sample group was drawn by word-of-mouth and informal advertisement of the study within the School of Computing at the university. A total of 23 respondents took part in this second pilot study. Demographics regarding the respondents were again not collected at this stage.

The data was once again transferred to an SQL server for analysis and links to the queries executed for the purpose of mining this data are included in Appendix I.

3.3.4 Results

The first mining query executed against the data set seeks to validate the overall usage of bimanual versus unimanual interaction to discover if the ratio of events maintains a consistent proportion to that of the first pilot study. The values remain very nearly identical with a marginal decrease in the usage of one object in each hand and a correspondingly marginal increase in single handed usage (Table 6).

Table 6 - Query E1B : Game type permutations for Pilot 2 study

Type	Incidents	As %
Bimanual	264	23.9
Unimanual	822	71.8
Duomanual	54	4.8

The pilot 2 studies averaged roughly 5 minutes of actual interaction time (actual time carrying an object) for each user by comparison to 2 minutes in the pilot 1 studies. An increase of an additional 133% attributable to the controlled study environment requiring consistent user interaction across all participants. Conversely the number of recorded actions dropped from 1757 to 1129 overall, a reduction of 35%. This is reflective of the more selective logging system increasing data relevancy.

Interaction methodology preference

Although the overall preference for unimanual interaction remains it is now possible to chart how that preference changes with repetition because all users have undertaken sixteen attempts at placing an object.

Fig. 25 shows the ratio of bimanual to unimanual interactions for each repetition and the computed trend-line demonstrates that users tend to have an increasing preference for bimanual interaction as the experiment progresses.

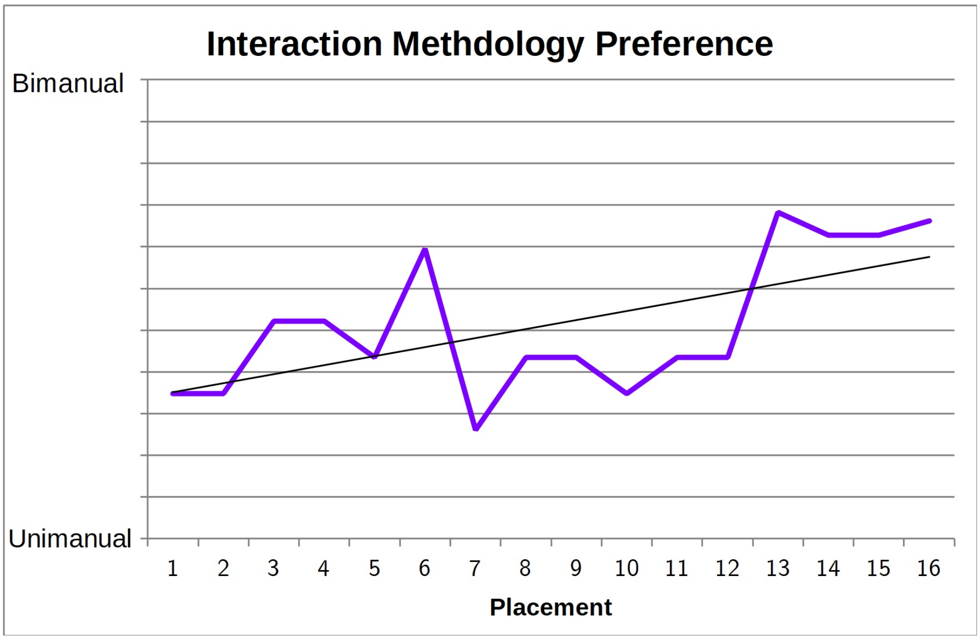

Fig. 25 - Interaction methodology trend

At the same time, overall placement accuracy also increases with repetition (Fig. 26).

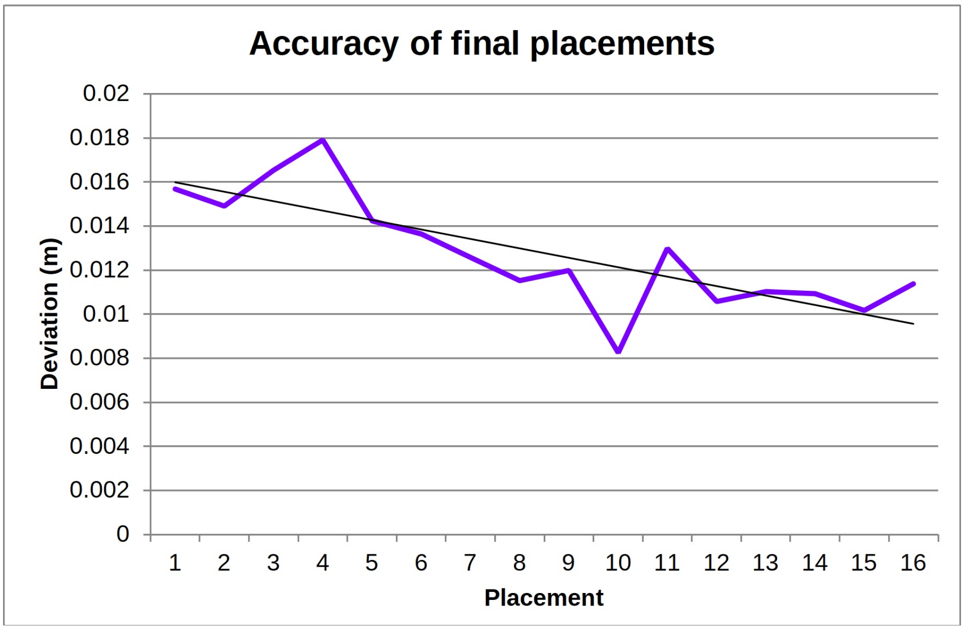

Fig. 26 - Placement precision trend

Although this confirms previous studies that suggest training is an important part of the process of any interaction system what is particularly interesting is how users seem to default so often to a unimanual interaction in virtual reality when it is arguably more natural to default to a bimanual one in real life. Is this a learned behaviour from having commonly utilised unimanual systems such as a mouse, or is it a weakness in this virtual reality system that leads users to perceive bimanual interaction as unnecessary or even perhaps lacking in benefit? Perhaps it is a function of the perception that objects will have no physical weight? Or perhaps a sense that the added precision of bimanual interaction is simply unnecessary? Or perhaps they simply feel they ought to be using unimanual interaction as the default state and are operating as they perceive is expected? Whilst this study does not address these questions it seems both surprising and interesting that, given a pair of identical controllers, portrayal of a pair of hands and bimanual instruction users nonetheless have a low initial attempt rate and low uptake rate for bimanual interaction.

Accuracy by game type

The lowest deviations from the optimal target occurred when assistive physics and indicators were both enabled (Fig. 27). This is continuing confirmation of numerous studies that suggest the use of assistive systems is beneficial. Interestingly, this does not hold true for the use of visual aids without physics simulation, perhaps because the indicator would give positive confirmation at any depth of penetration into the target object. It is possible that this positive bias caused the user to accept poorly positioned objects rather than trust their own visual judgement.

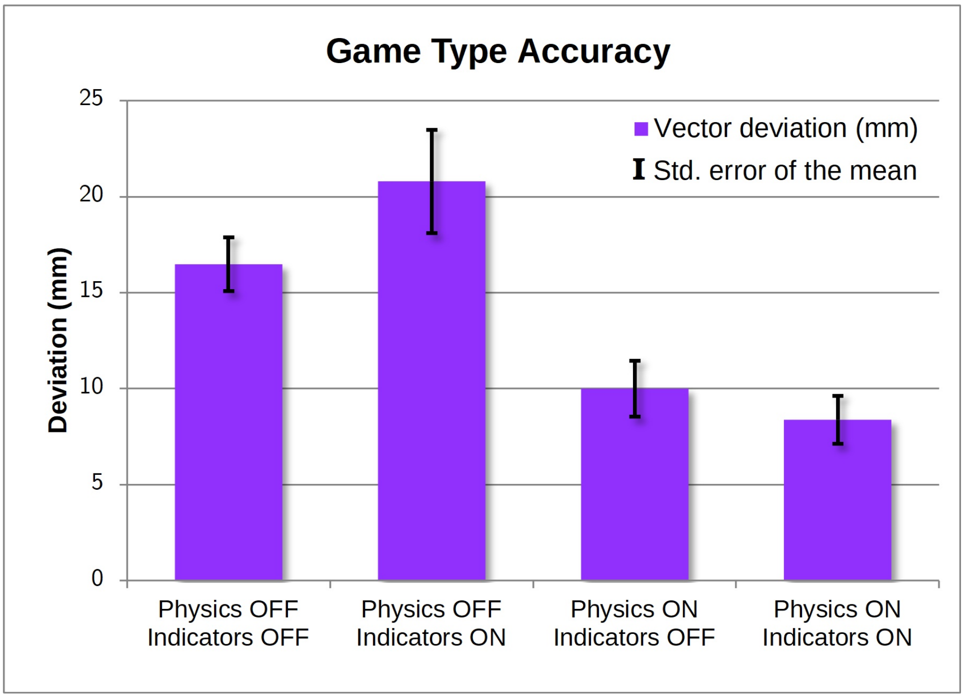

Fig. 27 - Accuracy of placement by game type

Isolating the data for deviations with physics disabled by axis it is clearly evident that the largest errors occur within the axis of object penetration in all modes of operation (Fig. 28). This upholds the theory of reliance upon the indicator system leading to excessive penetration since any position outside of the target, even if lateral and vertical precision were prefect, would still register negatively. Lateral error is the least in all modes whilst vertical error falls somewhere in-between and tends to be below the goal. It may be that angle of view of the user, which was universally above the target to varying degrees, may impact the precision of placement in some unexpected way.

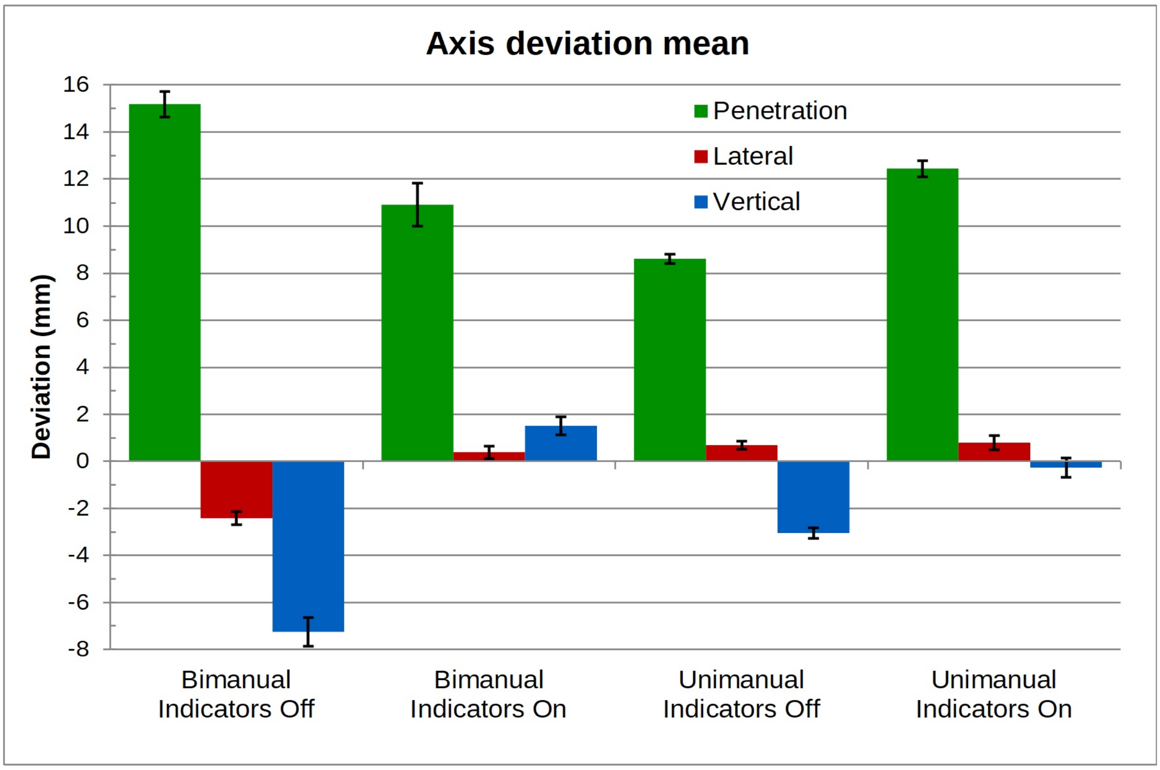

Fig. 28 - deviation mean by axis

Lack of indicators may affect vertical placement also since users tend to place below the object when it is disabled and very close to the correct position when enabled. It is possible that this is a result of depth perception limitations, a known constraint of many immersive systems. Generating a scatter diagram of the side-view placements (the lateral axis is fairly accurate throughout and can thus be ignored) and including both the region of confidence and approximate user gaze position does indeed appear to support this possibility (Fig. 29).

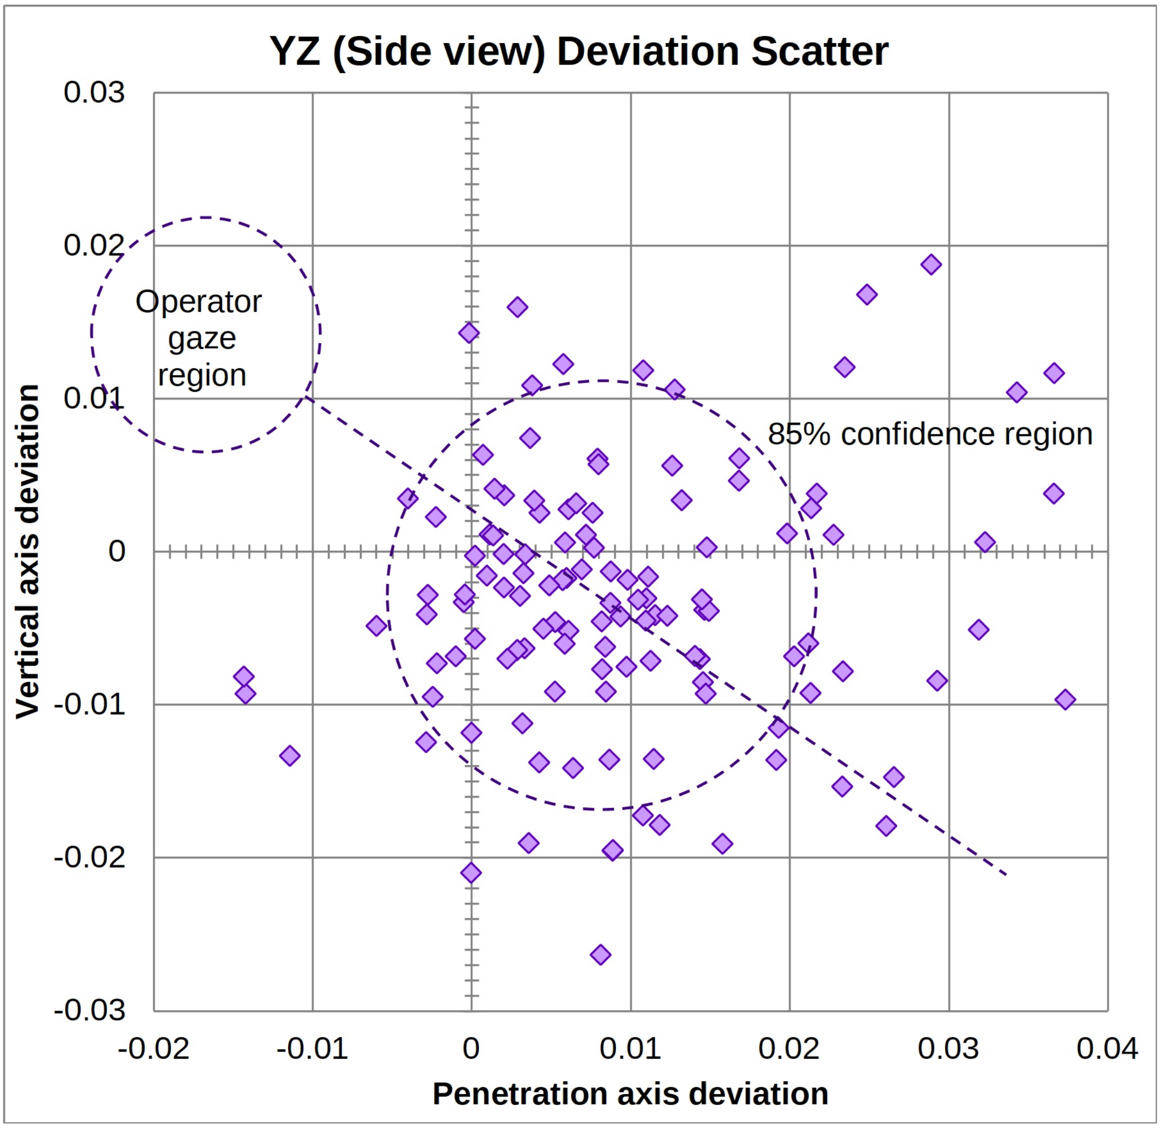

Fig. 29 - Scatter diagram of side-axis final placements

Taking the approximate operator gaze region, which was recorded for each data entry point, and including it with the region of confidence a direct line can be drawn through the centre of this region which correlates with a proportional user error some small distance below the ideal target and some slightly greater distance inside the target. If it was the case that the placement error is due to the limitations of depth perception in an immersive environment combined with positive affirmations of success from the indicator subsystem for any positive penetration depth, this could go some way to explaining this outcome.

Conclusion to the second pilot study

The data of the second study was more rigorously and selectively obtained. No errors in data collection process were found and user interaction was consistent and predictable throughout. The use of automation simplified the experiment execution whilst reducing the time required for each participant to complete the process.

The automated data gathering process eliminated the possibility of user error during the repetitive and intensive sanitisation process previously required, but upon reflection perhaps too little data was recorded. Whilst able to fulfil the objectives and answer the hypotheses presented initially it became evident that it was difficult to mine the data set in such a way as to demonstrate convincing proofs for unexpected behaviours. On balance, unless time is strictly limited, then too much data with a correspondingly difficult analysis process is a better compromise than too little data resulting in a limited ability to explore unexpected outcomes.

Having an indicator which offered positive affirmation for all depths of penetration resulted in a conflict of information between the indicator and the instruction to visually align the objects as precisely as possible. The outcome of this was the mean centre of placement was beyond the ideal target goal. This inadvertently created scenario does suggest that users may tend to pay attention to the indicator and factor it into their decision regarding satisfaction of placement.

Both the hypotheses initially presented are confirmed by the data. Users do tend to increasingly prefer bimanual interaction with repetition and users also tend to evidence higher accuracy of placement with assistive tools enabled than without. The caveat with this latter is that the assistive tools must not introduce conflict.

3.4 Final experiment design

3.4.1 Hypothesis

The primary study investigates methods of improving smoothness under bimanual operation. It does not compare unimanual to bimanual since it is largely accepted by the literature and borne out by the pilot studies that the latter is superior in most instances.

The core metric used for evaluating smoothness, the dependent variable, is the jerk metric. Few studies outside of the medical field have evaluated bimanual interaction for the purposes of smooth control and none were found which evaluated free-space bimanual interaction lacking kinematic constraint, demonstrated a definitive evaluation of indicator feedback impact on this, or recommended a preferred scheme for implementation.

The study proposes that the use of feedback indicators will have a significant and positive impact upon the smoothness of user performance. The primary hypothesis in its null and alternative forms is thus given as:

Ha_0 **Jerk levels will be improved through the use of indicators**

Ha_1 **Jerk levels will be inferior or equal without indicators**

This study proposes as its secondary hypothesis that visual feedback will have a more positive impact upon performance than haptic feedback. The secondary hypothesis in its null and alternative forms can thus be given as:

Hb_0 **Jerk levels will be improved with visual indicators**

Hb_1 **Jerk levels will be reduced with haptic indicators**

3.4.2 Design

There will be a single task that users repeat with variation. The task will require the user to carry an object from one location to another as smoothly as possible.

The task will be presented in one of three forms:

1. Without any feedback indicators
2. With visual feedback indicators
3. With haptic sensory feedback

Users will repeat each task four times in succession. This is for summative purposes in order to ensure evenly distributed data and provide the ability to reasonably exclude a single set should any spurious information be generated through user or system error. Four cycles is selected as being sufficient to provide reasonable tolerance for error elimination.

It is not intended to use this task cycle for the purpose of formative feedback. Although displaying a score or offering some other feedback between rounds to encourage user improvement across additional rounds could open up this possibility this would necessitate lengthening the duration of the user experience and risks impacting attention span. A quick, focussed evaluation is deemed of more merit in this instance.

This work is aimed specifically at professional operators who will be skilled in using some form of advanced interaction technology, possibly a teleoperated robot. Accordingly, casual and unskilled computer users are excluded through the screening process which permits only frequent users of a computer system, gaming console or other sufficiently advanced and comparable technology to continue to the experiment stage.

A within-subjects, repeated measures, crossover design is used whereby all users complete all tasks. This provides the advantage that each subject inherently serves as their own control and a smaller number of subjects can be involved for the same given level of confidence outcome. The risk, compared to a between-subjects design with regard to the impact of learning, is mitigated through the use of short experiment cycles. Carry-over should also be negligible and any residual effects is mitigated through counterbalancing.

The experiments are performed in counterbalanced order such that all subjects experience the same pattern of events but the data is is not confounded by a learning process across the twelve repetitions (Fig. 7). Counterbalancing is considered appropriate because there are no complex dependencies between the three forms and only the benefit of repetition and practice needs to be mitigated.

Table 7 - Counterbalancing sequence for the core study

Subject	None	Visual	Haptic
1	First	Second	Third
2	Third	First	Second
3	Second	Third	First

3.4.3 Configuration

In order to execute the experiment a room sufficiently large for the user to move freely in was required. The room was configured to have a space for mobility within the centre and an additional space for arms to extend into safely (Fig. 30).

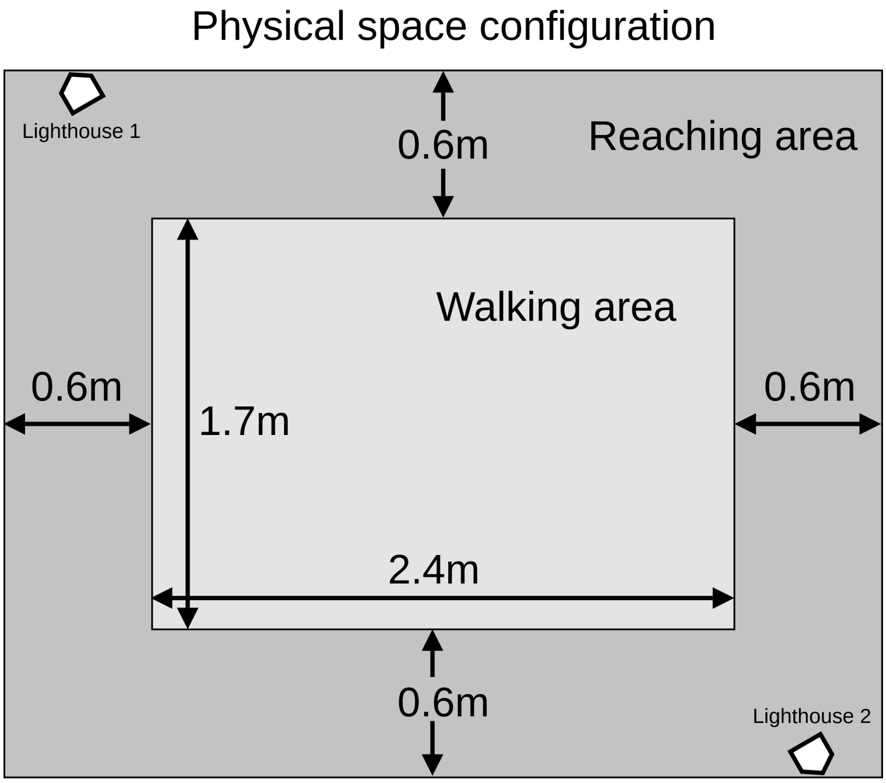

Fig. 30 - Room space diagram

The virtual gameplay area was mapped into this real-world space with visual cues so as to maximise available area whilst implying cues to safely limit user motion. Low level obstacles of tables and raised yellow gridwork were placed to discourage users from moving beyond the walking area (Fig. 31).

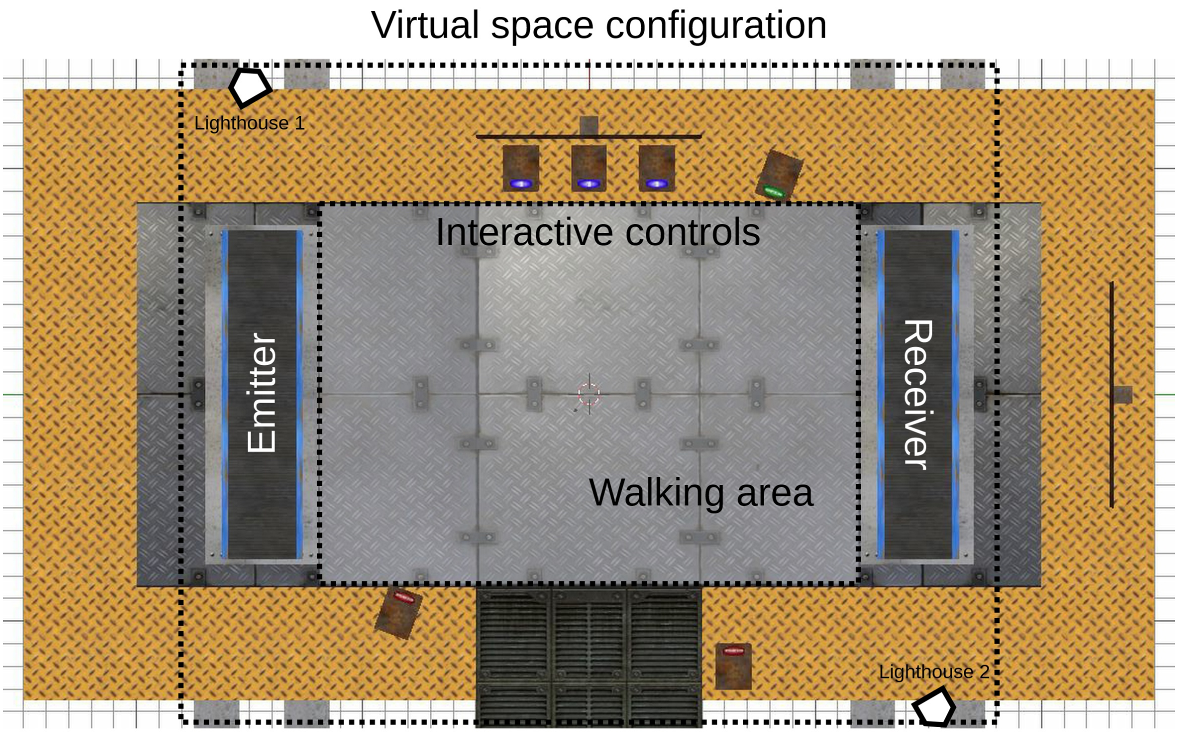

Fig. 31 - Simulation space diagram

The hardware used was a HTC Vive system with two laser lighthouses configured for synchronised Bluetooth mode operation. The computer was an AMD FX-8350 8 core CPU with 16GB of DDR3 RAM and a GeForce GTX1060 GPU. This configuration was tested with the developed software and was able to maintain the ideal 90Hz synchronised refresh rate for the experiment purposes to avoid drop-outs or the need for frame smoothing.

3.4.4 Software

Unreal Engine 4 predominantly uses a proprietary flow-chart based system called Blueprints (BP) for game development but this is supported by C++ code also. Blueprints is a new form of graphical programming designed to enable fast development and allow for quick experimental changes. The UE4 software comes with a quick-start template for the HTC Vive system that enables simple interaction and this is used as the basis for this work.

The Blueprints system is used to arrange the environment and to assign collision limits and interactive regions. The Blueprint system is also used to configure the C++ elements.

The most important change made to the UE4 system is that fast (predictive) tracker reporting is disabled. This default option allows for motion smoothing in games and other such consumer projects. Switching it off introduces some jitter which must be filtered out at the data processing stage and introduces the risk of jumps if the system is unable to maintain the refresh rate, but ensures that each positional sample recorded is synchronised exactly with the screen refresh and in-world representations of user interaction.

UE4 supports C++ in a specialised, non-standard form. Although all the capabilities of the C++ language are represented, in order to use this with BP it is necessary to allow UE4 to generate template files which include additional directives. These directives allow the BP system to generate code which connects the BP and C++ together. For practical purposes this means that C++ files should either inherit a pre-existing BP class and then add features to that, or they should be stand-alone tools, such as for mathematical computation.

The primary interaction object is a Blueprint Actor class extended with a Blueprint interface so that the system can provide positional data from the hand representations to which it maintains a reference. As an Actor class it also inherits the managed game loop implementation which allows for computation on each update cycle, known as a tick. The C++ class file is auto-generated as a template using UE4 and UE4 must be allowed to manage its constructors and linking. Such classes are prefixed with a capital 'A'.

UE4 allows Actor classes to utilise an interface through a managed template. This allows UE4 to perform reflection and code-linking even though C++ does not natively support this. Such interfaces must be generated by UE4 and are always prefixed with 'I'. Well-formed implementations must be suffixed with '`_Implementation`'.

The C++ portion of the UE4 simulation software utilises a total of five class files and two interfaces (Fig. 32). `ATestActor01` and `AGameButton` are the two actor classes which provide the user with objects to manipulate and buttons to control the game, respectively. Both of these classes utilise the `ISimpleGrab` interface which provides a means for the Blueprint system to pass controller information to them.

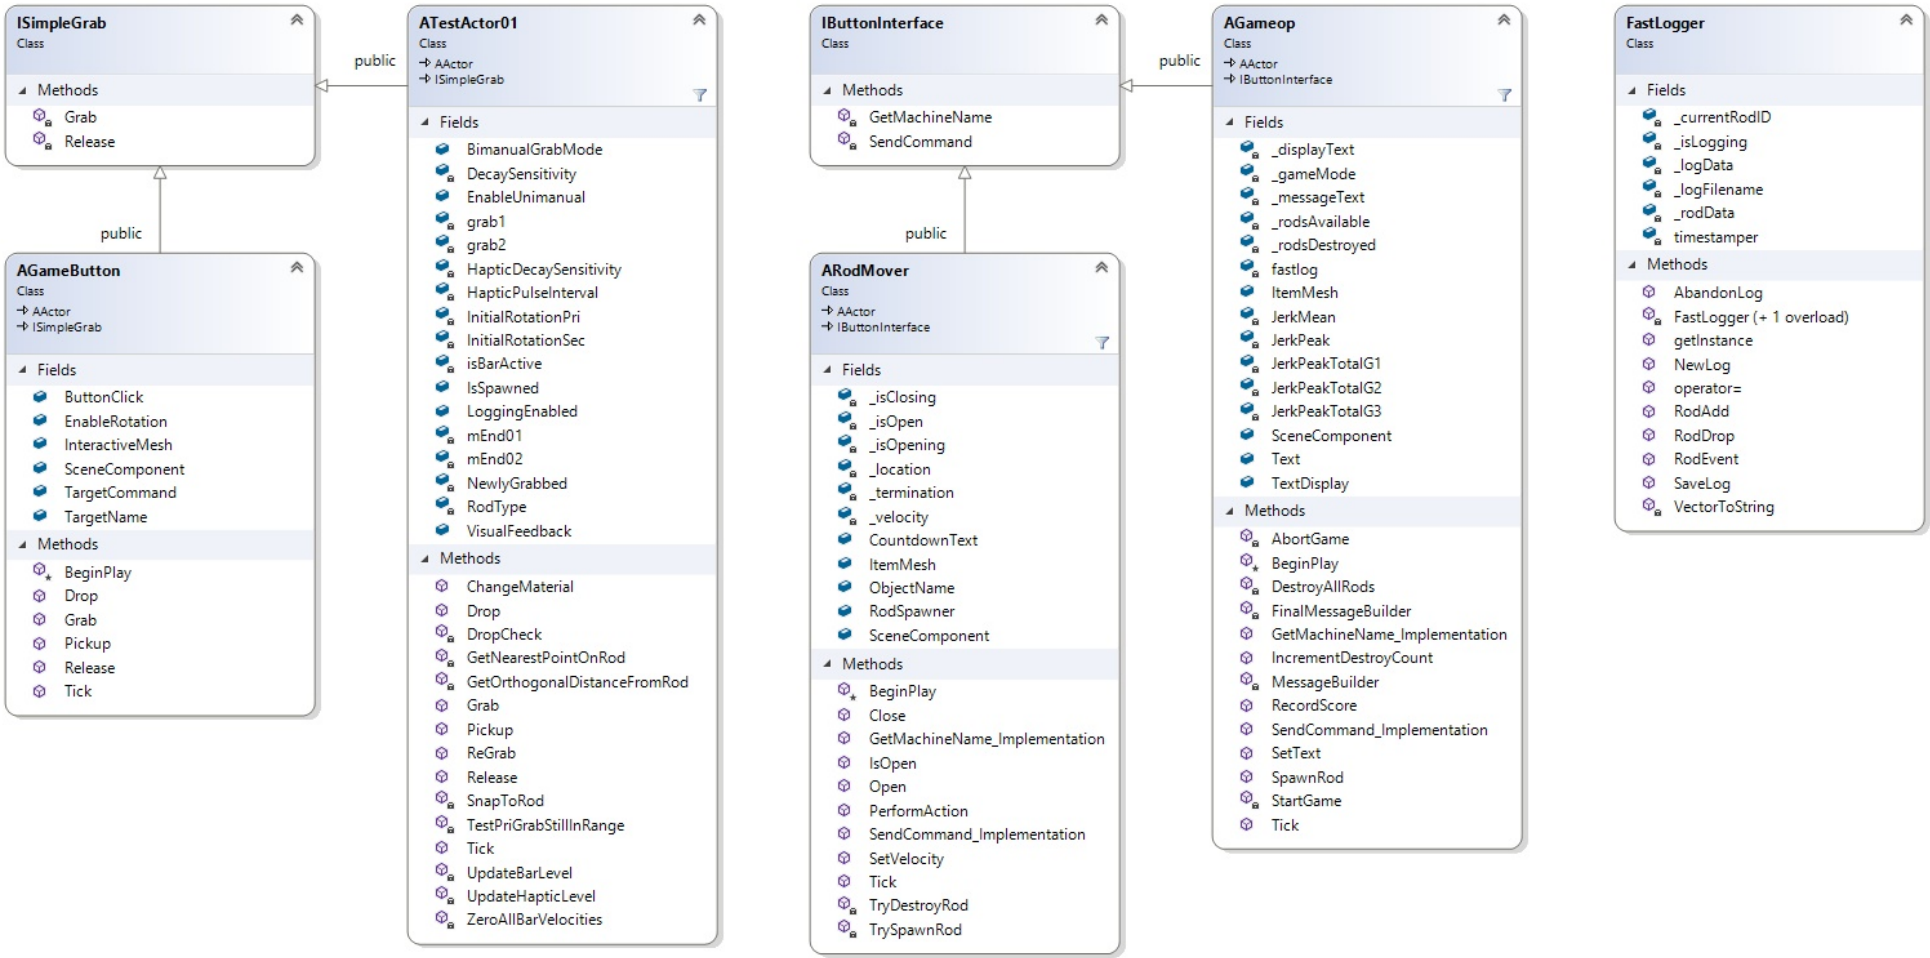

Fig. 32 - UE4 Class diagrams for C++ code

`AGameop` and `ARodMover` provide the game management tools and the system for spawning, destroying and animating the interactive objects, respectively. Both of these classes inherit from the `IButtonInterface` which allows the Blueprint based game system to broadcast commands to permit inter-component communication.

It should be noted that the `ATestActor01` class represents the primary game object with which the user interacts, henceforth referred to as a ‘Rod’ (Fig. 33). It is not possible to change class names in UE4 without risking severe compilation problems due to the way in which templating and reflection are implemented. This unfortunately ambiguous name is a hangover from early development experimentation before this limitation was understood.

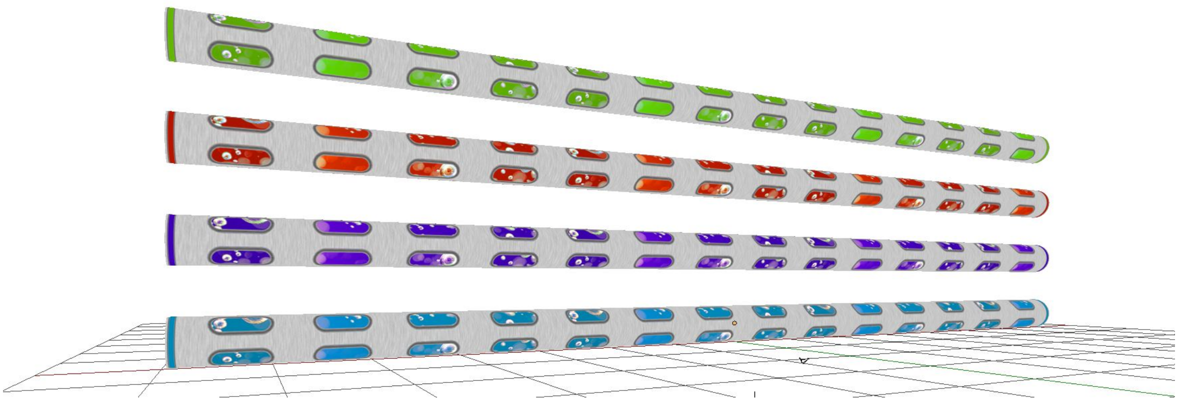

Fig. 33 - The primary interactive game objects aka ‘Rods’

Operation

The Blueprints system manages object collisions and on every game cycle creates a list of collision objects. If the user attempts an interaction the system checks the collision list to see if any utilises the `ISimpleButton` interface and, if so, calls that interface and passes a reference from the controller which called it.

If the interactive object is an `AGameButton` it broadcasts a preset message, set using the BP GUI, to all `IButtonInterfaces` and does not save the reference.

If the interactive object is an `ATestActor01` and the action is a grab then the reference is stored and a second reference is awaited. If the action is a release the reference is cleared. If the action is a second grab and an existing a reference exists then a second reference is stored, the rod is moved into the initial grab position and data begins to be sent to the logger system, which is instantiated as a singleton as necessary. Upon each cycle computation of rod position is performed as a centroid with rotation that passes through each garb position.

The `ARodMover` actor receives messages from the `IButtonInterfaces` and can act upon them to spawn a game object with specified characteristics, destroy a game object, animate an opening sequence or animate a closing sequence. This actor also loads and manages all associated models from the BP system that are required for this.

The `AGameop` class manages the game cycle and ensures that the correct type of objects are spawned and in the correct order. It also manages the information display to the user, orders the saving of logged data and resets the game area between games.

Finally, the `FastLogger` class is a singleton class that any `ARodMover` actor object can access to record data. It operates entirely in memory to ensure no slowdown and provides a tool for clearing the most recently buffered data if the user drops a rod out of bounds. The class also provides millisecond timing and file operations. This is the only class that is not part of a native UE4 BP template and thus lacks any prefix and uses purely standard C++.

3.4.5 Practicalities

Rod

The interactive object is designed to represent an item that implies a need for a bimanual interaction process. The length of 1.4m and diameter of 5cm was empirically selected to imply the requirement for a natural two handed grab, whilst the uniform simplicity of the object further implies no requirement for the user to grab it at any particular location. This is in contrast to the cube used in the pilot studies which was chosen to imply the option for unimanual or bimanual interaction was at user discretion.

The rod is made sufficiently long that it is unlikely any user will attempt to grab it at the extremes, this is by design to ensure that users always have the option of positioning and repositioning their hands within a preferred natural range without imposing limitation.

The cylindrical nature of the object ensures that rotation and twisting considerations can be discarded, both mathematically and from the perception of the user, and thus the research focus can be entirely upon the smoothness object motion as is sought.

Room

The room layout is created to require users to perform two process actions; carrying and twisting. Users must, at a minimum, walk some distance with the object and rotate whilst holding the object. This process creates a need for dynamic changes in motion. Although the operating area is constrained the overall room is made much larger to avoid a sense of claustrophobia whilst given plain walls to minimise distractions (Fig. 34).

Conveyor

The rods emerge from, and disappear into tanks fed from conveyors to ensure a sense of immersion where items vanishing and appearing might reduce that. The height of the conveyor is selected to be a natural height for most users so that all will operate without the need for crouching or stretching. This results in the centre of the rod being at 1.02m above floor level which is the approximate height at which utensils are used in a standard kitchen environment. The rod casings are designed to give the user a sense of space around the rod to encourage them to grab in any way they feel is correct.

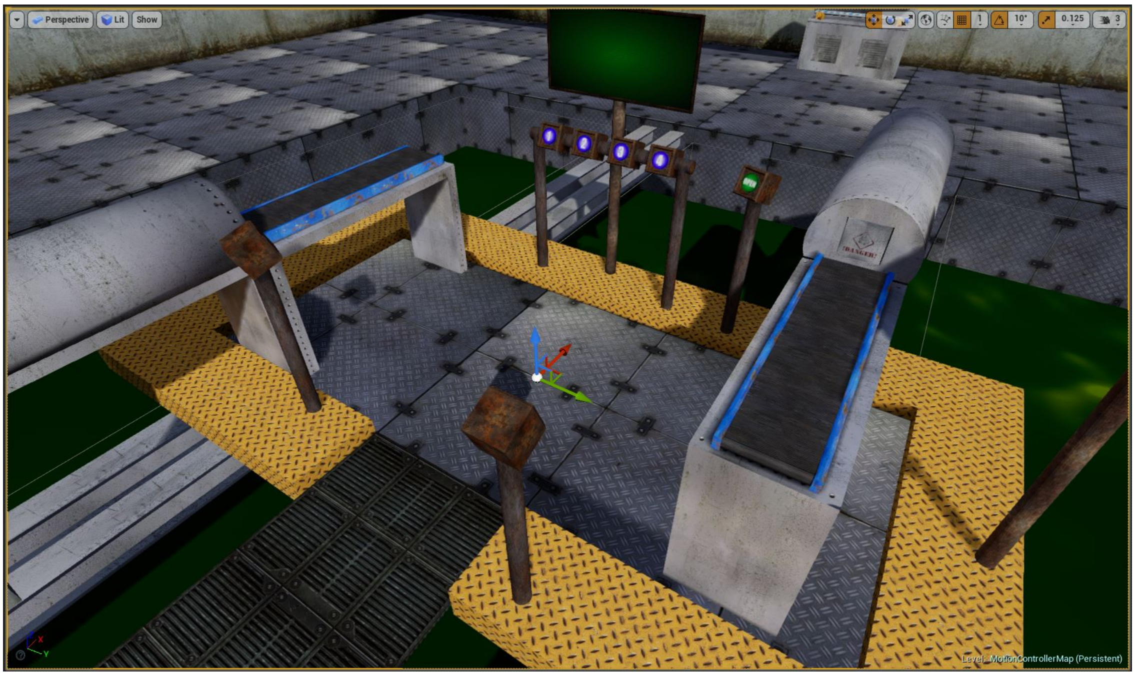

Fig. 34 - The game area including the additional space around it

3.4.6 Indicator feedback design

The visual and haptic indicators are designed to be as similar as possible in terms of function, response rate and sensitivity in order to minimise differences in evaluation. Both systems share a common jerk sensor system which evaluates the third derivative of position in real-time and uses a fixed scaling system to convert this into an integer score in the range of 0 to 8. Values above 8 are clipped at 8. The information is updated every frame. Frame rate is locked at 90 fps but the elapsed time is measure and used regardless to ensure the feedback is consistent should there be any fluctuation. The jerk data is taken from the far ends of the rod so is independent of hand position. The experiment supports different grab methodologies, such as shotgun, which locks the first grab in place and utilises the second for guiding the direction. The system evaluates each end independently and displays the highest jerk metric result.

The integer jerk metric is then passed to either visual or haptic indicators according to the rod type. Both visual and haptic indicators share a common constant which determines decay rate so that instantaneous fluctuations are presented clearly to the user. This was set, through empiric experimentation, to 100ms per segment or per rate-step for visual and haptic feedback correspondingly.

The visual feedback system was designed to float above the currently held rod and face the user at all times so as to ensure it remains visible (Fig. 35). It takes the form of an eight segment VU bar meter which are a common and easily understood paradigm.

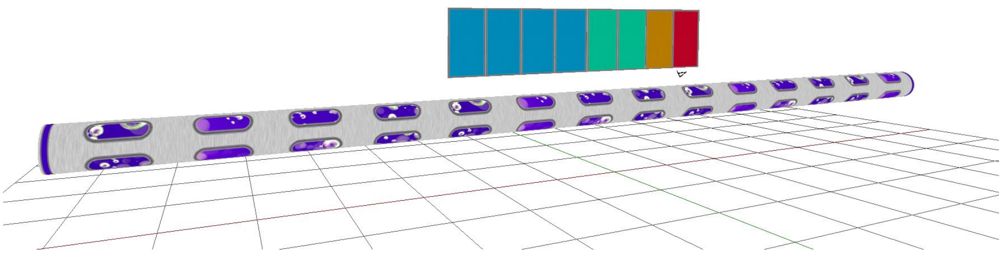

Fig. 35 - A rod with fully activated 8 bar VU meter

The haptic system uses pulse speed instead to present feedback. The rate increases across the 8 steps at exactly the same rate as does the visual meter. The haptic pulse rate is 2Hz at level 1 and 16Hz at level 8 with a stepped proportional increase between these values.

4. Data Mining

The raw three dimensional vector data for both hands and the far ends of the 1.4m rods is collected in synchronisation with the 90Hz refresh rate during the carry process. Each rod is separately identified by sequential number and type.

Table 8 - Rod IDs

ID	Indicator
VI	None
LA	Visual
GP	Haptic

The source data format are CSV files with a single file per participant. These files are merged into a single CSV dataset with the addition of a unique user ID number per file and a unique ID per row of data. A table is created upon the MS SQL server using the SQL server management tool and then the data is uploaded, verified and imported.

The mining process relies upon SQL queries involving WHILE loops and the LAG statement to enable insertions to be computed using recursive data. The use of these tools slows down the query processing time significantly. A query of the entire dataset to return fully computed average values across all tests took just under an hour using a commercial cloud-based database processing cluster.

The process can be examined in detail in the T-SQL query files provided (Appendix I) but is also briefly summarised here:

1. Compute Velocity in m/s for each axis $V_{(xyz)} = ([n]-[n_{-1}]) / \Delta t$
2. Smooth Velocity using a 20% filter $SV_{(xyz)} = ([SV_{-1}] + (([V]-[SV_{-1}]) * 0.2)$
3. Compute Acceleration m/s2 $A_{(xyz)} = ([SV]-[SV_{-1}]) / \Delta t$
4. Compute Jerk m/s2 $J_{(xyz)} = ([A]-[A_{-1}]) / \Delta t$
5. Compute jerk magnitude $J_{mag} = \sqrt{(J_x^2 + J_y^2 + J_z^2)}$

Each data file ID is paired with an associated questionnaire to enable queries to be arbitrarily designed when demographic patterns are sought. The demographic data is added to the jerk summary in an excel file along with per-rod jerk scoring to enable local queries without the need for re-running the SQL queries. This file is available via the links provided in Appendix I.

4.1 User demographics

A total of eleven participants took part in the final study.

Seven participants were female and four were male. Ages were grouped into bands. Five users matched the 22 to 34 years band and a further five were in the 35 to 44 years band. One user was in the 45 to 54 year band.

One user had limited vision in one eye and 36% wore glasses. One user was long-sighted but this did not affect their use of a HMD. One user wore contact lenses.

All users had at least frequent to advanced computer or gaming skills as per the screening requirements. 64% of users considered themselves to be advanced computer users, frequent gamers or both. 64% of users had previously used a HMD. 45% of users had previously experienced a CAVE or Octave environment.

54% of users were educated to undergraduate level or higher and 91% to A-level or higher.

With regard to post-evaluation user commentary all users reported that they experienced no nausea or motion sickness during the study. One user complained of discomfort but attributed this to the physical headset not the immersive experience. One user had prior experience of the PlayStation VR system and said the PSVR resulted in a sense of nausea where the HTC Vive used here did not. Another user stated that they preferred the sense of immersion with the HMD and felt it was superior to the CAVE type environment.

4.2 Data filtering and computation

The data is first filtered to eliminate the tracking noise. Any tracking system introduces variable errors from sample-to-sample but the relatively high sample rate of 90Hz and predictable random error fluctuations of less than 1 mm per sample makes this data set a candidate for the use of low-pass filtering.

A simple Infinite Impulse Response (IIR) single-pole filter in the form of an Exponential Averager is be implemented using the single-multiplier formula

$$y(n) = y(n-1) + \alpha[x(n) - y(n-1)]$$

where α (alpha) controls the cut-off frequency as a ratio of the sample-rate. The impact of this filter when applied to the raw data with a coefficient of 0.2 can be illustrated by extracting a small sample of velocity data and processing it (Fig.).

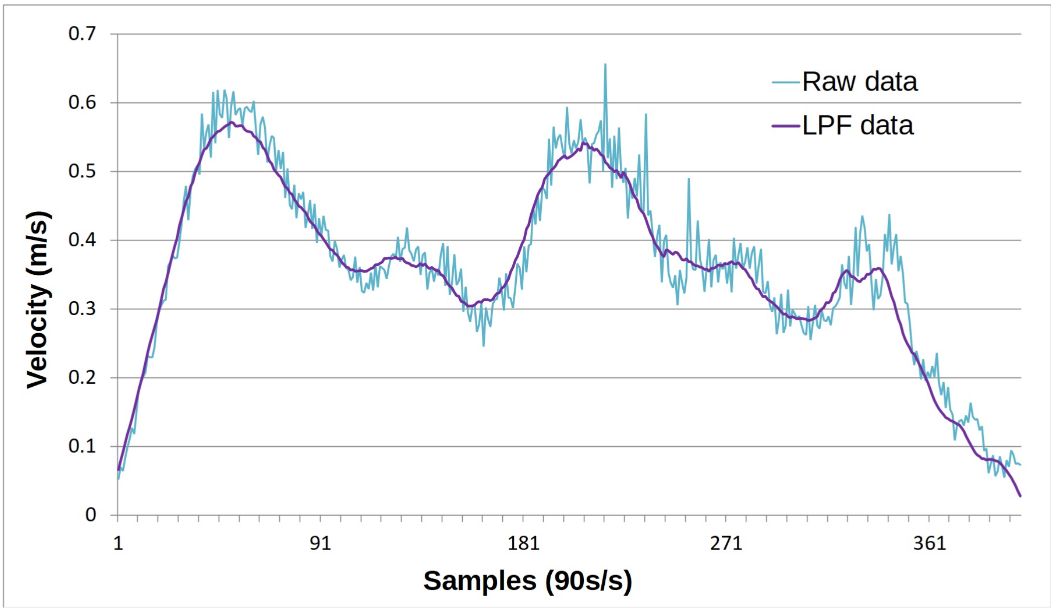

Fig. 36 - A randomly selected sample of x axis motion

The IIR filter is applied at the velocity processing level independently to each axis. Computation of acceleration and jerk are also performed discretely per axis. Final jerk is presented as a vector magnitude of the discrete jerk axes.

5. Analysis

The principal data sought is smoothness by feedback indicator type (FIT). The dependent variable is the mean jerk and the independent variable is the FIT. The analysis is quantitative. The study is a crossover study and all participants complete all trials. This specification makes this data suitable for analysis by parametric ANOVA but initially the data is more simply evaluated by FIT to find the mean and a Standard Error Margin calculated (Table 9). Each FIT result is the first-order filtered mean of 44 data sets of its respective type.

Table 9 - Average Jerk

FIT	Jerk	SEM	Time	SEM
None	83.3	3.5	10.8	0.7
Visual	63.3	3.0	18.1	1.0
Haptic	78.1	4.3	15.2	0.8

The low standard error margins suggests a high confidence that the use of either form of active feedback indicator resulted in reduced jerk and thus improved user performance. Visual indicators also outperformed haptic indicators by a visible margin (Fig. 9).

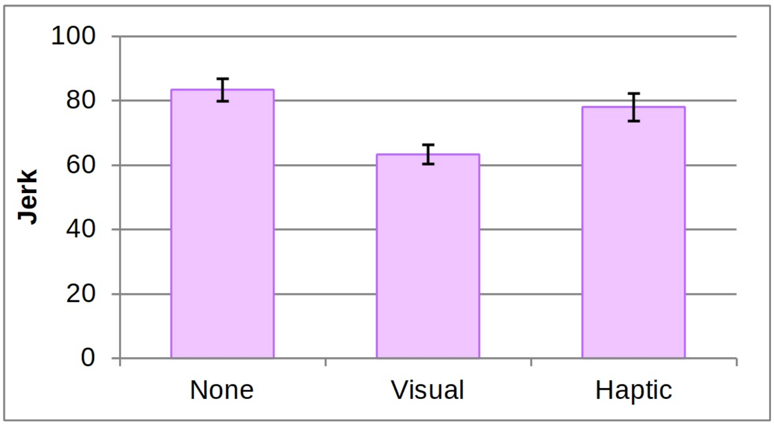

Fig. 37 - Comparing overall mean jerk by assistive indicator type

Notably, the increasing mean time taken to complete each task type is also congruent with reduced mean jerk / improved performance.

Progressing to a single factor ANOVA analysis (Table 10) the results show that the F factor is more than double the F critical and therefore the null hypothesis that indicator system do not make a difference to user performance is rejected (Table. 11). An α (alpha) level of 0.05 is generally accepted as a good threshold of statistical significance and the P-value returned is 0.0005 so therefore also supports the rejection of the null hypothesis.

Table 10 - Summary for test by FIT

Groups	Count	Sum	Average	Variance
None	44	3666.416	83.32764	535.9722
Visual	44	2787.143	63.34417	411.5023
Haptic	44	3435.678	78.08358	824.2157

Table 11 - ANOVA for test by FIT

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	9446.649	2	4723.325	7.998	0.000532	3.0663
Within Groups	76182.68	129	590.5634			
Total	**85629.33**	**131**				

Exploring the results beyond the hypothesis and splitting by gender suggests that female participants outperformed male participants as visualised in Fig. 38.

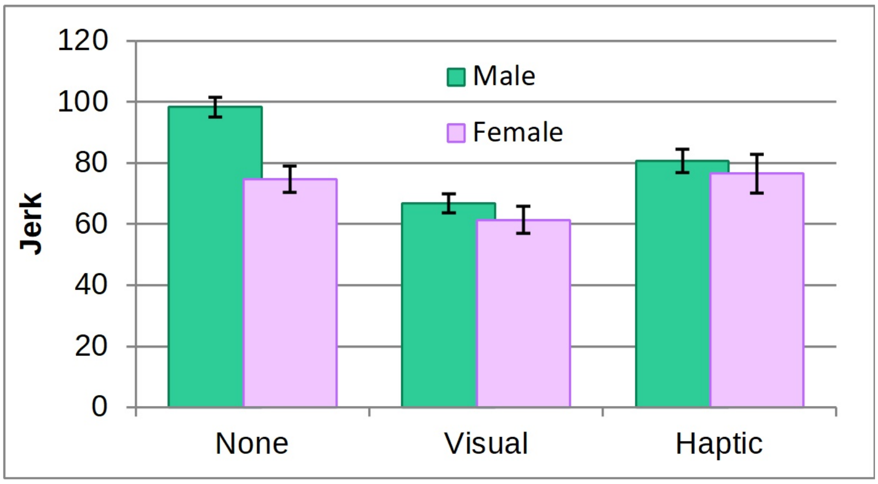

Fig. 38 - Mean jerk split by gender and by assistive indicator type displaying SEM bars

Performing a single factor ANOVA upon the aggregate results of male and female studies (Table 12) strongly supports this impromptu hypothesis, both in terms of F-factor and a significantly low P-value (Table 13).

Table 12 - Summary for test by gender

Groups	Count	Sum	Average	Variance
Female	84	5955.5	70.90	783.97
Male	48	3933.7	81.95	358.04

Table 13 - ANOVA for test by gender

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	3731.68	1	3731.69	5.92	0.0163	3.91
Within Groups	81897.64	130	630.98			
Total	**85629.33**	**131**				

Examining duration reveals that male users tend to complete the task much more quickly than female users (Table 14) and performing the respective ANOVA provides support for this with a high F value / F crit ratio and a very low P-value of 0.0003 (Table 15).

Table 14 - Mean duration by gender

FIT	Male (s)	Female (s)
None	7.5	12.6
Visual	14.9	19.9
Haptic	13.8	15.9

Table 15 - ANOVA for test by gender

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	512.56	1	512.56	13.78	0.0003	3.91
Within Groups	4836.95	130	37.21			
Total	**85629.33**	**131**				

However, the low population of the study overall and thus the possibility that these results are due to some other common but as-yet unidentified factor should be borne in mind. Accordingly, these additional findings are presented purely as an indicative interest piece and should be considered with some caution.

6. Conclusion

In a controlled experiment this work demonstrates that visual indicators have a more significant effect upon jerk performance than an equivalent haptic indicator and that either indicator is more significant than none. The ANOVAs executed displayed low P-values demonstrating high confidence in these results. Accordingly, the outcomes support both the primary and secondary hypotheses and clearly reject the corresponding null hypotheses.

The secondary demographic data suggests that females may generally yield lower jerk metrics than males and shows that females tend to spend more time completing the experiment. Even though the confidence level is reported as high a larger sample size would be preferred before claiming statistical significance.

As would be expected from a quantitative, computer science based study the most exciting data, in the most significant quantity, comes in a digital format and exploring this data will undoubtedly reveal further insights, especially in conjunction with the linked demographic responses. Accordingly however, the results presented here are brief and the data is instead made available via the links in Appendix I.

When this study began HMD systems such as the HTC Vive were not available. The few HMDs on offer were arguably limited in resolution, prohibitively expensive and required bespoke software, much like CAVEs did and, often, still do. Despite the perceived limitations of commercially available HMDs the impact is undeniable: The first two years of this work were spent developing a software framework for the Octave and it took a further year of work before an experiment could be run with any respectable rigour. Yet, when the Octave was unexpectedly out of service for five months it was possible to obtain a HMD, develop software for it, build an experiment and execute it in that time. The HMD ecosystem, the package of unified laser-tracked hardware and software has the potential to accelerate any research project an order of magnitude beyond a CAVE system.

Whilst CAVE systems have unique merits that are not easily challenged by HMD technology it seems often quietly bemoaned by researchers in their publications that a shortage of time restricted their works. Perhaps the HMD can finally unlock all that potential and free up researchers' time to perform better research, rather than be frustrated by having to develop underlying frameworks and technologies just to reach the start line.

6.1 Limitations and future work

This work does not include force-feedback style haptic systems and whilst comparing such devices to free-space devices might be of interest it would also have reduced the work to a device comparison. A parallel study designed solely around such a device with differing degrees of feedback versus visual indicators was out of scope in this instance.

The software provided (see Appendix I) is an improved version of this work that also logs individual hand positions in addition to the existing data. The core data logging remains identical to that used for this work however, and is interchangeable. The minor changes are the removal of white-space in the log files generated, automation of the experiment order cycling and a fix to clear the in game user display upon game start.

The data used for the final study was intended to only be a third pilot study but it was collected without errors or the need for significant changes and, as so often happens, time became limited so it became the core of this work and the improved software was unused.

It was observed in the initial study that users varied significantly in how they placed or moved their hands upon the rods and although not initially sought it was intended to also investigate this. It would also be desirable to investigate rotational jerk, which is rarely examined in any context, and the analysis of this is also supported.

Running this experiment again with greater numbers by using the provided software to investigate both hand placement and rotational jerk are the logical next steps.

Appendix I - Additional Materials

Most source materials including data, source and queries can be downloaded by visiting:

www.sm-robotics.com/phd.html

Additional materials are available by emailing the author at any of the following addresses:

sean@sm-robotics.com

sean@seanmandrake.com

sean@thefey.co.uk

Video materials can be most easily found by searching YouTube for the author's username:

seanmandrake

Appendix II - Ethical Approval

University of Salford
MANCHESTER

Research, Innovation and Academic Engagement Ethical Approval Panel

Research Centres Support Team
G0.3 Joule House
University of Salford
M5 4WT

T +44(0)161 295 5278

www.salford.ac.uk/

8 August 2016

Dear Sean,

RE: ETHICS APPLICATION ST16/108 – Improving Bimanual interaction within Virtual Worlds

Based on the information you provided, I am pleased to inform you that your application ST 16/108 has been approved.

If there are any changes to the project and/ or its methodology, please inform the Panel as soon as possible by contacting S&T-ResearchEthics@salford.ac.uk

Yours sincerely,

Prof Mohammed Arif
Chair of the Science & Technology Research Ethics Panel
Professor of Sustainability and Process Management,
School of Built Environment
University of Salford
Maxwell Building, The Crescent
Greater Manchester, UK M5 4WT
Phone: + 44 161 295 6829
Email: m.arif@salford.ac.uk

Bibliography

Achibet, M., Marchal, M., Argelaguet, F., Lecuyer, A., 2014. The Virtual Mitten: A novel interaction paradigm for visuo-haptic manipulation of objects using grip force, in: 2014 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2014 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, MN, USA, pp. 59–66. https://doi.org/10.1109/3DUI.2014.6798843

Anderson, L., Esser, J., Interrante, V., 2003. A virtual environment for conceptual design in architecture, in: Proceedings of the Workshop on Virtual Environments 2003 - EGVE '03. Presented at the the workshop, ACM Press, Zurich, Switzerland, pp. 57–63. https://doi.org/10.1145/769953.769960

Argelaguet, F., Andujar, C., 2013. A survey of 3D object selection techniques for virtual environments. Comput. Graph. 37, 121–136. https://doi.org/10.1016/j.cag.2012.12.003

Bachynskyi, M., Palmas, G., Oulasvirta, A., Weinkauf, T., 2015. Informing the Design of Novel Input Methods with Muscle Coactivation Clustering. ACM Trans. Comput.-Hum. Interact. 21, 1–25. https://doi.org/10.1145/2687921

Bacim, F., Kopper, R., Bowman, D.A., 2013. Design and evaluation of 3D selection techniques based on progressive refinement. Int. J. Hum.-Comput. Stud. 71, 785–802. https://doi.org/10.1016/j.ijhcs.2013.03.003

Balasubramanian, S., Melendez-Calderon, A., Burdet, E., 2012. A Robust and Sensitive Metric for Quantifying Movement Smoothness. IEEE Trans. Biomed. Eng. 59, 2126–2136. https://doi.org/10.1109/TBME.2011.2179545

Balasubramanian, S., Melendez-Calderon, A., Roby-Brami, A., Burdet, E., 2015. On the analysis of movement smoothness. J. NeuroEngineering Rehabil. 12. https://doi.org/10.1186/s12984-015-0090-9

Basdogan, C., Ho, C.-H., Srinivasan, M.A., Slater, M., 2000. An experimental study on the role of touch in shared virtual environments. ACM Trans. Comput.-Hum. Interact. 7, 443–460. https://doi.org/10.1145/365058.365082

Bau, O., Poupyrev, I., 2012. REVEL: tactile feedback technology for augmented reality. ACM Trans. Graph. 31, 1–11. https://doi.org/10.1145/2185520.2185585

Beattie, N., Horan, B., McKenzie, S., 2015. Taking the LEAP with the Oculus HMD and CAD - Plucking at thin Air? Procedia Technol. 20, 149–154. https://doi.org/10.1016/j.protcy.2015.07.025

Bertrand, J., Brickler, D., Babu, S., Madathil, K., Zelaya, M., Wang, T., Wagner, J., Gramopadhye, A., Luo, J., 2015. The role of dimensional symmetry on bimanual psychomotor skills education in immersive virtual environments, in: 2015 IEEE Virtual Reality (VR). Presented at the 2015 IEEE Virtual Reality (VR), IEEE, Arles, Camargue, Provence, France, pp. 3–10. https://doi.org/10.1109/VR.2015.7223317

Bogdan, N., Grossman, T., Fitzmaurice, G., 2014. HybridSpace: Integrating 3D freehand input and stereo viewing into traditional desktop applications, in: 2014 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2014 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, MN, USA, pp. 51–58. https://doi.org/10.1109/3DUI.2014.6798842

Bolas, M., 2014. Designing the user in user interfaces, in: Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology - UIST '14. Presented at the the 27th annual ACM symposium, ACM Press, Honolulu, Hawaii, USA, pp. 1–1. https://doi.org/10.1145/2642918.2642919

Bourdot, P., Convard, T., Picon, F., Ammi, M., Touraine, D., Vézien, J.-M., 2010. VR–CAD integration: Multimodal immersive interaction and advanced haptic paradigms for implicit edition of CAD models. Comput.-Aided Des. 42, 445–461. https://doi.org/10.1016/j.cad.2008.10.014

Boussemart, Y., Rioux, F., Rudzicz, F., Wozniewski, M., Cooperstock, J.R., 2004. A framework for 3D visualisation and manipulation in an immersive space using an untethered bimanual gestural interface, in: Proceedings of the ACM Symposium on Virtual Reality Software and Technology - VRST '04. Presented at the the ACM symposium, ACM Press, Hong Kong, p. 162. https://doi.org/10.1145/1077534.1077566

Bowman, D.A., Coquillart, S., Froehlich, B., Hirose, M., Kitamura, Y., Kiyokawa, K., Stuerzlinger, W., 2008. 3D User Interfaces: New Directions and Perspectives. IEEE Comput. Graph. Appl. 28, 20–36. https://doi.org/10.1109/MCG.2008.109

Bowman, D.A., Frohlich, B., Kitamura, Y., Stuerzlinger, W., 2005. New Directions in 3D User Interfaces, in: IEEE Virtual Reality Conference 2005 (VR'05). Presented at the IEEE Virtual Reality Conference 2005 (VR'05), IEEE, Bonn, Germany, pp. 312–312. https://doi.org/10.1109/VR.2005.58

Bowman, D.A., Johnson, D.B., Hodges, L.F., 1999. Testbed Evaluation of Virtual Environment Interaction Techniques, in: Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST '99. ACM, New York, NY, USA, pp. 26–33. https://doi.org/10.1145/323663.323667

Bowman, D.A., Kruijff, E., LaViola, J.J., Poupyrev, I., 2001. An Introduction to 3-D User Interface Design. Presence Teleoperators Virtual Environ. 10, 96–108. https://doi.org/10.1162/105474601750182342

Bowman, D.A., McMahan, R.P., Ragan, E.D., 2012. Questioning naturalism in 3D user interfaces. Commun. ACM 55, 78. https://doi.org/10.1145/2330667.2330687

Bowman, D.A., Ray, A.A., Gutierrez, M.S., Mauldon, M., Dove, J.E., Westman, E., Setareh, M., 2006. Engineering in Three Dimensions: Immersive Virtual Environments, Interactivity, and 3D User Interfaces for Engineering Applications, in: GeoCongress 2006. Presented at the GeoCongress 2006, American Society of Civil Engineers, Atlanta, Georgia, United States, pp. 1–17. https://doi.org/10.1061/40803(187)6

Bridgwater, L.B., Ihrke, C.A., Diftler, M.A., Abdallah, M.E., Radford, N.A., Rogers, J.M., Yayathi, S., Askew, R.S., Linn, D.M., 2012. The Robonaut 2 hand - designed to do work with tools, in: 2012 IEEE International Conference on Robotics and Automation. Presented at the 2012 IEEE International Conference on Robotics and Automation, pp. 3425–3430. https://doi.org/10.1109/ICRA.2012.6224772

Bruder, G., Sanz, F.A., Olivier, A.-H., Lecuyer, A., 2015. Distance estimation in large immersive projection systems, revisited, in: 2015 IEEE Virtual Reality (VR). Presented at the 2015 IEEE Virtual Reality (VR), IEEE, Arles, Camargue, Provence, France, pp. 27–32. https://doi.org/10.1109/VR.2015.7223320

Bruder, G., Steinicke, F., Sturzlinger, W., 2013. Effects of visual conflicts on 3D selection task performance in stereoscopic display environments, in: 2013 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2013 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Orlando, FL, pp. 115–118. https://doi.org/10.1109/3DUI.2013.6550207

Buxton, W., Myers, B., 1986. A Study in Two-handed Input, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI ’86. ACM, New York, NY, USA, pp. 321–326. https://doi.org/10.1145/22627.22390

Capobianco, A., Veit, M., Bechmann, D., 2009. A Preliminary Study of Two-Handed Manipulation for Spatial Input Tasks in a 3D Modeling Application, in: Lopez Jaquero, V., Montero Simarro, F., Molina Masso, J.P., Vanderdonckt, J. (Eds.), Computer-Aided Design of User Interfaces VI. Springer London, London, pp. 77–88. https://doi.org/10.1007/978-1-84882-206-1_8

Chen, J., Bowman, D.A., 2009. Domain-Specific Design of 3D Interaction Techniques: An Approach for Designing Useful Virtual Environment Applications. Presence Teleoperators Virtual Environ. 18, 370–386. https://doi.org/10.1162/pres.18.5.370

Cho, I., Wartell, Z., 2015. Evaluation of a bimanual simultaneous 7DOF interaction technique in virtual environments, in: 2015 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2015 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Arles, France, pp. 133–136. https://doi.org/10.1109/3DUI.2015.7131738

Cutler, L.D., Frölich, B., Hanrahan, P., 1997. Two-handed direct manipulation on the responsive workbench, in: Proceedings of the 1997 Symposium on Interactive 3D Graphics - SI3D ’97. Presented at the the 1997 symposium, ACM Press, Providence, Rhode Island, United States, p. 107-ff. https://doi.org/10.1145/253284.253315

De Araùjo, B.R., Casiez, G., Jorge, J.A., 2012. Mockup Builder: Direct 3D Modeling on and Above the Surface in a Continuous Interaction Space, in: Proceedings of Graphics Interface 2012, GI ’12. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, pp. 173–180.

DeFanti, T.A., Dawe, G., Sandin, D.J., Schulze, J.P., Otto, P., Girado, J., Kuester, F., Smarr, L., Rao, R., 2009. The StarCAVE, a third-generation CAVE and virtual reality OptIPortal. Future Gener. Comput. Syst. 25, 169–178. https://doi.org/10.1016/j.future.2008.07.015

Diftler, M.A., Mehling, J.S., Abdallah, M.E., Radford, N.A., Bridgwater, L.B., Sanders, A.M., Askew, R.S., Linn, D.M., Yamokoski, J.D., Permenter, F.A., Hargrave, B.K., Platt, R., Savely, R.T., Ambrose, R.O., 2011. Robonaut 2 - The first humanoid robot in space, in: 2011 IEEE International Conference on Robotics and Automation. Presented at the 2011 IEEE International Conference on Robotics and Automation (ICRA), IEEE, Shanghai, China, pp. 2178–2183. https://doi.org/10.1109/ICRA.2011.5979830

Duval, T., Lecuyer, A., Thomas, S., 2006. SkeweR: a 3D Interaction Technique for 2-User Collaborative Manipulation of Objects in Virtual Environments, in: 3D User Interfaces (3DUI’06). Presented at the 3D User Interfaces (3DUI’06), pp. 69–72. https://doi.org/10.1109/VR.2006.119

Estrada, S., O’Malley, M.K., Duran, C., Schulz, D., Bismuth, J., 2014. On the development of objective metrics for surgical skills evaluation based on tool motion, in: 2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC). Presented at the 2014 IEEE International Conference on Systems, Man and Cybernetics - SMC, IEEE, San Diego, CA, USA, pp. 3144–3149. https://doi.org/10.1109/SMC.2014.6974411

Feng, J., Cho, I., Wartell, Z., 2015. Comparison of Device-Based, One and Two-Handed 7DOF Manipulation Techniques, in: Proceedings of the 3rd ACM Symposium on Spatial User Interaction - SUI '15. Presented at the the 3rd ACM Symposium, ACM Press, Los Angeles, California, USA, pp. 2–9. https://doi.org/10.1145/2788940.2788942

Fiorentino, M., Uva, A.E., Dellisanti Fabiano, M., Monno, G., 2010. Improving bi-manual 3D input in CAD modelling by part rotation optimisation. Comput.-Aided Des. 42, 462–470. https://doi.org/10.1016/j.cad.2008.12.002

Flores-Abad, A., Ma, O., Pham, K., Ulrich, S., 2014. A review of space robotics technologies for on-orbit servicing. Prog. Aerosp. Sci. 68, 1–26. https://doi.org/10.1016/j.paerosci.2014.03.002

Fong, T., Provencher, C., Micire, M., Diftler, M., Berka, R., Bluethmann, B., Mittman, D., 2012. The Human Exploration Telerobotics project: Objectives, approach, and testing, in: 2012 IEEE Aerospace Conference. Presented at the 2012 IEEE Aerospace Conference, pp. 1–9. https://doi.org/10.1109/AERO.2012.6187043

Froehlich, B., Hochstrate, J., Skuk, V., Huckauf, A., 2006. The GlobeFish and the GlobeMouse: two new six degree of freedom input devices for graphics applications, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '06. Presented at the the SIGCHI conference, ACM Press, Montréal, Québec, Canada, p. 191. https://doi.org/10.1145/1124772.1124802

Garcia-Robledo, P., Ortego, J., Barrio, J., Galiana, I., Ferre, M., Aracil, R., 2009. Multifinger haptic interface for bimanual manipulation of virtual objects, in: 2009 IEEE International Workshop on Haptic Audio Visual Environments and Games. Presented at the 2009 IEEE International Workshop on Haptic Audio visual Environments and Games (HAVE 2009), IEEE, Lecco, Italy, pp. 30–35. https://doi.org/10.1109/HAVE.2009.5356128

Giachritsis, C., Barrio, J., Ferre, M., Wing, A., Ortego, J., 2009. Evaluation of weight perception during unimanual and bimanual manipulation of virtual objects, in: World Haptics 2009 - Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. Presented at the 2009 world Haptics Conference (WHC 2009), IEEE, Salt Lake City, UT, USA, pp. 629–634. https://doi.org/10.1109/WHC.2009.4810836

Graham, J.L., Manuel, S.G., Johannes, M.S., Armiger, R.S., 2011. Development of a multi-modal haptic feedback system for dexterous robotic telemanipulation, in: 2011 IEEE International Conference on Systems, Man, and Cybernetics. Presented at the 2011 IEEE International Conference on Systems, Man and Cybernetics - SMC, IEEE, Anchorage, AK, USA, pp. 3548–3553. https://doi.org/10.1109/ICSMC.2011.6084219

Guiard, Y., 1987. Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model. J. Mot. Behav. 19, 486–517. https://doi.org/10.1080/00222895.1987.10735426

Handelman, D.A., Franken, G.H., Komsuoglu, H., 2010. Agile and dexterous robot for inspection and EOD operations, in: Gerhart, G.R., Gage, D.W., Shoemaker, C.M. (Eds.), . Presented at the SPIE Defense, Security, and Sensing, Orlando, Florida, p. 769211. https://doi.org/10.1117/12.851251

Harrison, B., Oehmen, R., Robertson, A., Robertson, B., De Cruz, P., Khan, R., Fick, D., 2017. Through the eye of the master: The use of Virtual Reality in the teaching of surgical hand preparation, in: 2017 IEEE 5th International Conference on Serious Games and

Applications for Health (SeGAH). Presented at the 2017 IEEE 5th International Conference on Serious Games and Applications for Health (SeGAH), IEEE, Perth, Australia, pp. 1–6. https://doi.org/10.1109/SeGAH.2017.7939269

Hinckley, K., Pausch, R., Proffitt, D., Kassell, N.F., 1998. Two-handed virtual manipulation. ACM Trans. Comput.-Hum. Interact. 5, 260–302. https://doi.org/10.1145/292834.292849

Hinckley, K., Pausch, R., Proffitt, D., Patten, J., Kassell, N., 1997. Cooperative bimanual action, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI ’97. Presented at the the SIGCHI conference, ACM Press, Atlanta, Georgia, United States, pp. 27–34. https://doi.org/10.1145/258549.258571

Hinton, M.A., Johannes, M.S., Zeher, M.J., Kozlowski, M.V., 2011. IMPLEMENTING A COMMON ARCHITECTURE FOR EOD ROBOTIC SYSTEMS 13.

Hogan, N., Sternad, D., 2009. Sensitivity of Smoothness Measures to Movement Duration, Amplitude, and Arrests. J. Mot. Behav. 41, 529–534. https://doi.org/10.3200/35-09-004-RC

Hughes, C.E., Lelin Zhang, Schulze, J.P., Edelstein, E., Macagno, E., 2013. CaveCAD: Architectural design in the CAVE, in: 2013 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2013 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Orlando, FL, pp. 193–194. https://doi.org/10.1109/3DUI.2013.6550244

Israel, J.H., Mauderli, L., Greslin, L., 2013. Mastering digital materiality in immersive modelling, in: Proceedings of the International Symposium on Sketch-Based Interfaces and Modeling - SBIM ’13. Presented at the the International Symposium, ACM Press, Anaheim, California, p. 15. https://doi.org/10.1145/2487381.2487388

Jenkins, O.C., 2008. Sparse control for high-DOF assistive robots. Intell. Serv. Robot. 1, 135–141. https://doi.org/10.1007/s11370-007-0013-0

Jerald, J., Mlyniec, P., Yoganandan, A., Rubin, A., Paullus, D., Solotko, S., 2013. MakeVR: A 3D world-building interface, in: 2013 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2013 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Orlando, FL, pp. 197–198. https://doi.org/10.1109/3DUI.2013.6550246

Jiao, X., Deng, H., Wang, F., 2010. An Investigation of Two-Handed Manipulation and Related Techniques in Multi-touch Interaction, in: 2010 International Conference on Machine Vision and Human-Machine Interface. Presented at the 2010 International Conference on Machine Vision and Human-machine Interface, IEEE, Kaifeng, China, pp. 565–568. https://doi.org/10.1109/MVHI.2010.27

Khare, R., Bascom, R., Higgins, W.E., 2015. Hands-Free System for Bronchoscopy Planning and Guidance. IEEE Trans. Biomed. Eng. 62, 2794–2811. https://doi.org/10.1109/TBME.2015.2401514

Kil, Y.J., Renzulli, P., Kreylos, O., Hamann, B., Monno, G., Staadt, O., 2005. 3D warp brush: interactive free-form modeling on the responsive workbench, in: IEEE Proceedings. VR 2005. Virtual Reality, 2005. Presented at the IEEE Proceedings. VR 2005. Virtual Reality, 2005., pp. 279–280. https://doi.org/10.1109/VR.2005.1492794

Kron, A., Schmidt, G., Petzold, B., Zah, M.I., Hinterseer, P., Steinbach, E., 2004. Disposal of explosive ordnances by use of a bimanual haptic telepresence system, in: IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA ’04. 2004. Presented at the IEEE International Conference on Robotics and Automation, 2004.

Proceedings. ICRA '04. 2004, IEEE, New Orleans, LA, USA, pp. 1968-1973 Vol.2. https://doi.org/10.1109/ROBOT.2004.1308112

Kunert, A., Kulik, A., Huckauf, A., Fröhlich, B., 2007. A Comparison of Tracking- and Controller-based Input for Complex Bimanual Interaction in Virtual Environments, in: Proceedings of the 13th Eurographics Conference on Virtual Environments, EGVE'07. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, pp. 43–52. https://doi.org/10.2312/EGVE/IPT_EGVE2007/043-052

Lampton, D.R., Knerr, B.W., Goldberg, S.L., Bliss, J.P., Moshell, J.M., Blau, B.S., 1994. The Virtual Environment Performance Assessment Battery Vepab: Development and Evaluation1. Presence Teleoper Virtual Env. 3, 145–157. https://doi.org/10.1162/pres.1994.3.2.145

Leitner, J., Luciw, M., Forster, A., Schmidhuber, J., 2014. Teleoperation of a 7 DOF Humanoid Robot Arm Using Human Arm Accelerations and EMG Signals 8.

Li, J., Cho, I., Wartell, Z., 2015. Evaluation of 3D virtual cursor offset techniques for navigation tasks in a multi-display virtual environment, in: 2015 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2015 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Arles, France, pp. 59–66. https://doi.org/10.1109/3DUI.2015.7131727

Liu, L., Miyake, S., Maruyama, N., Akahane, K., Sato, M., 2014. Development of Two-Handed Multi-finger Haptic Interface SPIDAR-10, in: Auvray, M., Duriez, C. (Eds.), Haptics: Neuroscience, Devices, Modeling, and Applications. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 176–183. https://doi.org/10.1007/978-3-662-44196-1_22

Lubos, P., Bruder, G., Steinicke, F., 2014. Analysis of direct selection in head-mounted display environments, in: 2014 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2014 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, MN, USA, pp. 11–18. https://doi.org/10.1109/3DUI.2014.6798834

Lucas, J.F., Bowman, D.A., Chen, J., Wingrave, C.A., 2005. Design and Evaluation of 3D Multiple Object Selection Techniques. Unpublished 1.

Meaney, D.F., 2000. Dynamic Mechanical Stretch of Organotypic Brain Slice Cultures Induces Differential Genomic Expression: Relationship to Mechanical Parameters. J. Biomech. Eng. 122, 224. https://doi.org/10.1115/1.429650

Meyrueis, V., Paljic, A., Fuchs, P., 2009. D^3: an immersive aided design deformation method, in: Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology - VRST '09. Presented at the the 16th ACM Symposium, ACM Press, Kyoto, Japan, p. 179. https://doi.org/10.1145/1643928.1643968

Mine, M., 2003. Towards Virtual Reality for the Masses: 10 Years of Research at Disney's VR Studio, in: Proceedings of the Workshop on Virtual Environments 2003, EGVE '03. ACM, New York, NY, USA, pp. 11–17. https://doi.org/10.1145/769953.769955

Mine, M., Yoganandan, A., Coffey, D., 2014. Making VR work: building a real-world immersive modeling application in the virtual world, in: Proceedings of the 2nd ACM Symposium on Spatial User Interaction - SUI '14. Presented at the the 2nd ACM symposium, ACM Press, Honolulu, Hawaii, USA, pp. 80–89. https://doi.org/10.1145/2659766.2659780

Mine, M.R., Brooks, F.P., Sequin, C.H., 1997. Moving objects in space: exploiting proprioception in virtual-environment interaction, in: Proceedings of the 24th Annual Conference on

Computer Graphics and Interactive Techniques - SIGGRAPH ’97. Presented at the the 24th annual conference, ACM Press, Not Known, pp. 19–26. https://doi.org/10.1145/258734.258747

Ministry Of Defence, 2018. British Army receives pioneering bomb disposal robots [WWW Document]. GOV.UK. URL https://www.gov.uk/government/news/british-army-receives-pioneering-bomb-disposal-robots (accessed 1.19.19).

Mohamadipanah, H., Parthiban, C., Law, K., Nathwani, J., Maulson, L., DiMarco, S., Pugh, C., 2016. Hand smoothness in laparoscopic surgery correlates to psychomotor skills in virtual reality, in: 2016 IEEE 13th International Conference on Wearable and Implantable Body Sensor Networks (BSN). Presented at the 2016 IEEE 13th International Conference on Wearable and Implantable Body Sensor Networks (BSN), IEEE, San Francisco, CA, USA, pp. 242–246. https://doi.org/10.1109/BSN.2016.7516267

Murakami, E.A.Y., Shibata, K., Zheng, X.-Z., Ito, K., 2000. Human control characteristics in bilateral micro-teleoperation system, in: 2000 26th Annual Conference of the IEEE Industrial Electronics Society. IECON 2000. 2000 IEEE International Conference on Industrial Electronics, Control and Instrumentation. 21st Century Technologies. Presented at the 2000 26th Annual Conference of the IEEE Industrial Electronics Society. IECON 2000. 2000 IEEE International Conference on Industrial Electronics, Control and Instrumentation. 21st Century Technologies, pp. 602–607 vol.1. https://doi.org/10.1109/IECON.2000.973218

Murayama, J., Bougrila, L., Luo, Y., Akahane, K., Hasegawa, S., Hirsbrunner, B., Sato, M., 2004. SPIDAR G&G: A Two-Handed Haptic Interface for Bimanual VR Interaction 9.

Murray, N., 2011. Contextual Interaction Support in 3D Worlds, in: 2011 IEEE/ACM 15th International Symposium on Distributed Simulation and Real Time Applications. Presented at the 2011 IEEE/ACM 15th International Symposium on Distributed Simulation and Real Time Applications (DS-RT), IEEE, Salford, United Kingdom, pp. 58–63. https://doi.org/10.1109/DS-RT.2011.19

Nguyen, A., Banic, A., 2015. 3DTouch: A Wearable 3D Input Device for 3D Applications. IEEE Virtual Real. VR 7.

Niinimäki, M., Tahiroglu, K., 2012. AHNE: a novel interface for spatial interaction, in: Proceedings of the 2012 ACM Annual Conference Extended Abstracts on Human Factors in Computing Systems Extended Abstracts - CHI EA ’12. Presented at the the 2012 ACM annual conference extended abstracts, ACM Press, Austin, Texas, USA, p. 1031. https://doi.org/10.1145/2212776.2212378

O’Brien, T.M., Keefe, D.F., Laidlaw, D.H., 2008. A Case Study in Using Gestures and Bimanual Interaction to Extend a high-DOF Input Device, in: Proceedings of the 2008 Symposium on Interactive 3D Graphics and Games, I3D ’08. ACM, New York, NY, USA, pp. 9:1–9:1. https://doi.org/10.1145/1342250.1357021

O’Hara, K., Dastur, N., Carrell, T., Gonzalez, G., Sellen, A., Penney, G., Varnavas, A., Mentis, H., Criminisi, A., Corish, R., Rouncefield, M., 2014. Touchless interaction in surgery. Commun. ACM 57, 70–77. https://doi.org/10.1145/2541883.2541899

Oikonomidis, I., Kyriazis, N., Argyros, A.A., 2012. Tracking the articulated motion of two strongly interacting hands, in: 2012 IEEE Conference on Computer Vision and Pattern Recognition. Presented at the 2012 IEEE Conference on Computer Vision and Pattern Recognition

(CVPR), IEEE, Providence, RI, pp. 1862–1869. https://doi.org/10.1109/CVPR.2012.6247885

Paljic, A., Coquillart, S., Burkhardt, J.-M., Richard, P., 2002. A Study of Distance of Manipulation on the Responsive WorkbenchTM, in: IEEE VR Conference 2002. Presented at the Immersive Projection Symposia IPT, IEEE, p. 8.

Pfeiffer, M., Stuerzlinger, W., 2015. 3D virtual hand pointing with EMS and vibration feedback, in: 2015 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2015 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Arles, France, pp. 117–120. https://doi.org/10.1109/3DUI.2015.7131735

Pinho, M.S., Bowman, D.A., Freitas, C.M.D.S., 2002. Cooperative Object Manipulation in Immersive Virtual Environments: Framework and Techniques, in: Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST ’02. ACM, New York, NY, USA, pp. 171–178. https://doi.org/10.1145/585740.585769

Pollard, N.S., Hodgins, J.K., Riley, M.J., Atkeson, C.G., 2002. Adapting human motion for the control of a humanoid robot, in: Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292). Presented at the 2002 IEEE International Conference on Robotics and Automation, IEEE, Washington, DC, USA, pp. 1390–1397. https://doi.org/10.1109/ROBOT.2002.1014737

Polys, N.F., Kim, S., Bowman, D.A., 2007. Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments. Comput. Animat. Virtual Worlds 18, 19–38. https://doi.org/10.1002/cav.159

Ponto, K., Tredinnick, R., Bartholomew, A., Roy, C., Szafir, D., Greenheck, D., Kohlmann, J., 2013. SculptUp: A rapid, immersive 3D modeling environment, in: 2013 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2013 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Orlando, FL, pp. 199–200. https://doi.org/10.1109/3DUI.2013.6550247

Poupyrev, I., Ichikawa, T., Weghorst, S., Billinghurst, M., 1998. Egocentric Object Manipulation in Virtual Environments: Empirical Evaluation of Interaction Techniques. Comput. Graph. Forum 17, 41–52. https://doi.org/10.1111/1467-8659.00252

Poupyrev, I., Weghorst, S., Billinghurst, M., Ichikawa, T., 1997. A Framework and Testbed for Studying Manipulation Techniques for Immersive VR, in: Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST ’97. ACM, New York, NY, USA, pp. 21–28. https://doi.org/10.1145/261135.261141

Prachyabrued, M., Borst, C.W., 2014. Visual feedback for virtual grasping, in: 2014 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2014 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, MN, USA, pp. 19–26. https://doi.org/10.1109/3DUI.2014.6798835

Rahimian, F.P., Ibrahim, R., 2011. Impacts of VR 3D sketching on novice designers’ spatial cognition in collaborative conceptual architectural design. Des. Stud. 32, 255–291. https://doi.org/10.1016/j.destud.2010.10.003

Rioux, F., Rudzicz, F., Wozniewski, M., 2004. The Modellers’ Apprentice - Gesture-Based 3D Design in Immersive Environments. Unpublished 6.

Ruddle, R.A., 2002. Symmetric and Asymmetric Action Integration During Cooperative Object Manipulation in Virtual Environments 29.

Scerbo, S., Bowman, D., 2012. Design issues when using commodity gaming devices for virtual object manipulation, in: Proceedings of the International Conference on the Foundations of Digital Games - FDG '12. Presented at the the International Conference, ACM Press, Raleigh, North Carolina, p. 294. https://doi.org/10.1145/2282338.2282406

Schemali, L., Eisemann, E., 2014. Design and evaluation of mouse cursors in a stereoscopic desktop environment, in: 2014 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2014 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, MN, USA, pp. 67–70. https://doi.org/10.1109/3DUI.2014.6798844

Schultheis, U., Jerald, J., Toledo, F., Yoganandan, A., Mlyniec, P., 2012. Comparison of a two-handed interface to a wand interface and a mouse interface for fundamental 3D tasks, in: 2012 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2012 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Costa Mesa, CA, pp. 117–124. https://doi.org/10.1109/3DUI.2012.6184195

Seagull, F.J., Miller, P., George, I., Mlyniec, P., Park, A., 2009. Interacting in 3D Space: Comparison of a 3D Two-handed Interface to a Keyboard-and-mouse Interface for Medical 3D Image Manipulation. Hum. Factors Ergon. Soc. Annu. Meet. Proc. 53, 2004–2008. https://doi.org/10.1518/107118109X12524444845234

Song, P., Goh, W.B., Hutama, W., Fu, C.-W., Liu, X., 2012. A handle bar metaphor for virtual object manipulation with mid-air interaction, in: Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems - CHI '12. Presented at the the 2012 ACM annual conference, ACM Press, Austin, Texas, USA, p. 1297. https://doi.org/10.1145/2207676.2208585

Stannus, S., Fu, W.-T., Lucieer, A., 2014. Natural 7DoF input for 3D navigation, in: Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures the Future of Design - OzCHI '14. Presented at the the 26th Australian Computer-Human Interaction Conference, ACM Press, Sydney, New South Wales, Australia, pp. 216–219. https://doi.org/10.1145/2686612.2686646

Stenholt, R., 2012. Efficient selection of multiple objects on a large scale, in: Proceedings of the 18th ACM Symposium on Virtual Reality Software and Technology - VRST '12. Presented at the the 18th ACM symposium, ACM Press, Toronto, Ontario, Canada, p. 105. https://doi.org/10.1145/2407336.2407357

Stoll, E., Letschnik, J., Walter, U., Artigas, J., Kremer, P., Preusche, C., Hirzinger, G., 2009. On-orbit servicing. IEEE Robot. Autom. Mag. 16, 29–33. https://doi.org/10.1109/MRA.2009.934819

Stoll, E., Letschnik, J., Wilde, M., Saenz-Otero, A., Varatharajoo, R., Artigas, J., 2012. The future role of relay satellites for orbital telerobotics. Adv. Space Res. 50, 864–880. https://doi.org/10.1016/j.asr.2012.05.014

Takala, T.M., 2014. RUIS: a toolkit for developing virtual reality applications with spatial interaction, in: Proceedings of the 2nd ACM Symposium on Spatial User Interaction - SUI '14. Presented at the the 2nd ACM symposium, ACM Press, Honolulu, Hawaii, USA, pp. 94–103. https://doi.org/10.1145/2659766.2659774

Takala, T.M., Makarainen, M., Hamalainen, P., 2013. Immersive 3D modeling with Blender and off-the-shelf hardware, in: 2013 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2013 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Orlando, FL, pp. 191–192. https://doi.org/10.1109/3DUI.2013.6550243

Teather, R.J., Stuerzlinger, W., 2014. Visual aids in 3D point selection experiments, in: Proceedings of the 2nd ACM Symposium on Spatial User Interaction - SUI '14. Presented at the the 2nd ACM symposium, ACM Press, Honolulu, Hawaii, USA, pp. 127–136. https://doi.org/10.1145/2659766.2659770

Thomas, B.H., 2012. A survey of visual, mixed, and augmented reality gaming. Comput. Entertain. 10, 1–33. https://doi.org/10.1145/2381876.2381879

Tunstel Jr., E.W., Wolfe, K.C., Kutzer, M.D.M., Johannes, M.S., Brown, C.Y., Katyal, K.D., Para, M.P., Zeher, M.J., 2013. Recent Enhancements to Mobile Bimanual Robotic Teleoperation with Insight Toward Improving Operator Control. JOHNS HOPKINS APL Tech. Dig. 32, 11.

Ulinski, A., Wartell, Z., Hodges, L.F., 2007. Bimanual task division preferences for volume selection, in: Proceedings of the 2007 ACM Symposium on Virtual Reality Software and Technology - VRST '07. Presented at the the 2007 ACM symposium, ACM Press, Newport Beach, California, p. 217. https://doi.org/10.1145/1315184.1315228

van Rhijn, A., Mulder, J.D., 2006. Spatial input device structure and bimanual object manipulation in virtual environments, in: Proceedings of the ACM Symposium on Virtual Reality Software and Technology - VRST '06. Presented at the the ACM symposium, ACM Press, Limassol, Cyprus, p. 51. https://doi.org/10.1145/1180495.1180507

Veit, M., Capobianco, A., Bechmann, D., 2008. Consequence of Two-handed Manipulation on Speed, Precision and Perception on Spatial Input Task in 3D Modelling Applications. Unpublished 14.

Wang, J., Leach, O., Lindeman, R.W., 2013. DIY World Builder: An immersive level-editing system, in: 2013 IEEE Symposium on 3D User Interfaces (3DUI). Presented at the 2013 IEEE Symposium on 3D User Interfaces (3DUI), IEEE, Orlando, FL, pp. 195–196. https://doi.org/10.1109/3DUI.2013.6550245

Ware, C., Jessome, D.R., 1988. Using the bat: a six-dimensional mouse for object placement. IEEE Comput. Graph. Appl. 8, 65–70. https://doi.org/10.1109/38.20319

Weller, R., Zachmann, G., 2012. User Performance in complex bi-manual haptic manipulation with 3 DOFs vs. 6 DOFs, in: 2012 IEEE Haptics Symposium (HAPTICS). Presented at the 2012 IEEE Haptics Symposium (HAPTICS), IEEE, Vancouver, BC, Canada, pp. 315–322. https://doi.org/10.1109/HAPTIC.2012.6183808

Wolfe, K.C., Kutzer, M.D.M., Tunstel, E.W., 2016. Exploiting redundancy to improve bimanual telepresent manipulation, in: 2016 World Automation Congress (WAC). Presented at the 2016 World Automation Congress (WAC), pp. 1–6. https://doi.org/10.1109/WAC.2016.7583020

Wright, T., de Ribaupierre, S., Eagleson, R., 2017. Design and evaluation of an augmented reality simulator using leap motion. Healthc. Technol. Lett. 4, 210–215. https://doi.org/10.1049/htl.2017.0070

Yang, R.S., Strozzi, A.G., Lau, A., Lutteroth, C., Chan, Y.H., Delmas, P., 2012. Bimanual Natural User Interaction for 3D Modelling Application Using Stereo Computer Vision, in: Proceedings of the 13th International Conference of the NZ Chapter of the ACM's Special Interest Group on Human-Computer Interaction, CHINZ '12. ACM, New York, NY, USA, pp. 44–51. https://doi.org/10.1145/2379256.2379264

Yin, J., Liu, H., 2010. The empirical study of a novel bimanual scrolling technique, in: The 2010 IEEE International Conference on Information and Automation. Presented at the 2010 International Conference on Information and Automation (ICIA), IEEE, Harbin, China, pp. 110–113. https://doi.org/10.1109/ICINFA.2010.5512346

Zahedi, E., Rahmat-Khah, H., Dargahi, J., Zadeh, M., 2017. Virtual reality based training: Evaluation of user performance by capturing upper limb motion, in: 2017 IEEE Virtual Reality (VR). Presented at the 2017 IEEE Virtual Reality (VR), IEEE, Los Angeles, CA, USA, pp. 255–256. https://doi.org/10.1109/VR.2017.7892273

Zhai, S., Milgram, P., 1998. Quantifying coordination in multiple DOF movement and its application to evaluating 6 DOF input devices, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '98. Presented at the the SIGCHI conference, ACM Press, Los Angeles, California, United States, pp. 320–327. https://doi.org/10.1145/274644.274689