

Bangor University

DOCTOR OF PHILOSOPHY

An Advanced Virtual Environment for Rugby Skills Training

Miles, Helen

Award date:
2014

Awarding institution:
Bangor University

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Declaration and Consent

Details of the Work

I hereby agree to deposit the following item in the digital repository maintained by Bangor University and/or in any other repository authorized for use by Bangor University.

Author Name:

Title:

Supervisor/Department:

Funding body (if any):

Qualification/Degree obtained:

This item is a product of my own research endeavours and is covered by the agreement below in which the item is referred to as “the Work”. It is identical in content to that deposited in the Library, subject to point 4 below.

Non-exclusive Rights

Rights granted to the digital repository through this agreement are entirely non-exclusive. I am free to publish the Work in its present version or future versions elsewhere.

I agree that Bangor University may electronically store, copy or translate the Work to any approved medium or format for the purpose of future preservation and accessibility. Bangor University is not under any obligation to reproduce or display the Work in the same formats or resolutions in which it was originally deposited.

Bangor University Digital Repository

I understand that work deposited in the digital repository will be accessible to a wide variety of people and institutions, including automated agents and search engines via the World Wide Web.

I understand that once the Work is deposited, the item and its metadata may be incorporated into public access catalogues or services, national databases of electronic theses and dissertations such as the British Library’s EThOS or any service provided by the National Library of Wales.

I understand that the Work may be made available via the National Library of Wales Online Electronic Theses Service under the declared terms and conditions of use (<http://www.llgc.org.uk/index.php?id=4676>). I agree that as part of this service the National Library of Wales may electronically store, copy or convert the Work to any approved medium or format for the purpose of future preservation and accessibility. The National Library of Wales is not under any obligation to reproduce or display the Work in the same formats or resolutions in which it was originally deposited.

Statement 1:

This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree unless as agreed by the University for approved dual awards.

Signed (candidate)

Date

Statement 2:

This thesis is the result of my own investigations, except where otherwise stated. Where correction services have been used, the extent and nature of the correction is clearly marked in a footnote(s).

All other sources are acknowledged by footnotes and/or a bibliography.

Signed (candidate)

Date

Statement 3:

I hereby give consent for my thesis, if accepted, to be available for photocopying, for inter-library loan and for electronic storage (subject to any constraints as defined in statement 4), and for the title and summary to be made available to outside organisations.

Signed (candidate)

Date

NB: Candidates on whose behalf a bar on access has been approved by the Academic Registry should use the following version of **Statement 3:**

Statement 3 (bar):

I hereby give consent for my thesis, if accepted, to be available for photocopying, for inter-library loans and for electronic storage (subject to any constraints as defined in statement 4), after expiry of a bar on access.

Signed (candidate)

Date

Statement 4:

Choose **one** of the following options

a)	I agree to deposit an electronic copy of my thesis (the Work) in the Bangor University (BU) Institutional Digital Repository, the British Library ETHOS system, and/or in any other repository authorized for use by Bangor University and where necessary have gained the required permissions for the use of third party material.	
b)	I agree to deposit an electronic copy of my thesis (the Work) in the Bangor University (BU) Institutional Digital Repository, the British Library ETHOS system, and/or in any other repository authorized for use by Bangor University when the approved bar on access has been lifted.	
c)	I agree to submit my thesis (the Work) electronically via Bangor University's e-submission system, however I opt-out of the electronic deposit to the Bangor University (BU) Institutional Digital Repository, the British Library ETHOS system, and/or in any other repository authorized for use by Bangor University, due to lack of permissions for use of third party material.	

Options B should only be used if a bar on access has been approved by the University.

In addition to the above I also agree to the following:

1. That I am the author or have the authority of the author(s) to make this agreement and do hereby give Bangor University the right to make available the Work in the way described above.
2. That the electronic copy of the Work deposited in the digital repository and covered by this agreement, is identical in content to the paper copy of the Work deposited in the Bangor University Library, subject to point 4 below.
3. That I have exercised reasonable care to ensure that the Work is original and, to the best of my knowledge, does not breach any laws – including those relating to defamation, libel and copyright.
4. That I have, in instances where the intellectual property of other authors or copyright holders is included in the Work, and where appropriate, gained explicit permission for the inclusion of that material in the Work, and in the electronic form of the Work as accessed through the open access digital repository, *or* that I have identified and removed that material for which adequate and appropriate permission has not been obtained and which will be inaccessible via the digital repository.
5. That Bangor University does not hold any obligation to take legal action on behalf of the Depositor, or other rights holders, in the event of a breach of intellectual property rights, or any other right, in the material deposited.
6. That I will indemnify and keep indemnified Bangor University and the National Library of Wales from and against any loss, liability, claim or damage, including without limitation any related legal fees and court costs (on a full indemnity bases), related to any breach by myself of any term of this agreement.

Signature: Date :

**An Advanced
Virtual Environment
for Rugby
Skills Training**



PRIFYSGOL
BANGOR
UNIVERSITY

Thesis submitted in accordance with the requirements of
the Bangor University for the degree of Doctor in Philosophy

by

Helen C. Miles

January 2014

ABSTRACT

There is growing interest in utilising virtual environments (VEs) in the context of sports. In particular there is a desire to be able to improve sensorimotor skills rather than just using a VE as a tool for strategy analysis, or entertainment. While there have been a number of VEs developed for the sports of tennis, football and baseball, very little work has been done for the game of rugby. The main aim of this thesis is to address this gap in the research. Passing is chosen as the skill in question, as it is considered a core, fundamental skill of the game that is sometimes forgotten amongst more complex training regimes. No previous work has been undertaken to build a virtual environment to train passing skills for rugby, and so a prototype system has been designed to explore it's potential.

The system (VERST: a Virtual Environment for Rugby Skills Training) was designed to have a user facing a large screen with a virtual scene containing targets; the user holds a real rugby ball (tethered to prevent damage to equipment) and must throw it at the virtual targets.

The system was tested in two exploratory experiments featuring a throwing task and a verbal estimation task. 10 participants were recruited to assess the perception of depth for the virtual targets and the suitability of the design of the system. Different configurations of the hardware were tested: the participant's position relative to the screen, the use of stereoscopy and the use of a floor screen. The results suggests that the difference in intended target distances is being correctly perceived, but that the subjects

were not throwing that distance. Though the reason for this is not clear, it is possible that the virtual scene chosen for the task lacked sufficient depth cues. It is also possible that the task was not appropriately designed, as the task of throwing the ball forward is an illegal move in a game of rugby. Despite negative results from the throwing task, VERST has proven to be a useful platform for investigating the optimal set up of a virtual environment for training ball passing skills in rugby.

Contents

Abstract	i
List of Figures	x
List of Tables	xviii
Acknowledgements	xx
1 Introduction	1
1.1 Context and Motivation	1
1.2 Hypothesis	2
1.3 Contributions	2
1.3.1 List of Publications	2
1.4 Thesis Structure	4
2 Background	5
2.1 Background: First Concepts	6
2.1.1 Introduction	6
2.1.2 Motor Skill Learning	6
2.1.3 Virtual Environments	7
2.1.4 Virtual Environments for Sports Training	8
2.1.5 Validation and Verification	8
2.1.6 Sports-themed Computer Games	13

2.1.7	Modeling the Physical Forces	15
2.1.7.1	The Principles of Motion	15
2.1.7.2	Physics Engines	18
2.1.7.3	Modeling the Motion of Sports Projectiles	19
2.1.8	Components Required for a Sports Virtual Environment	20
2.2	Hardware and Software Used for Sports Virtual Environments	21
2.2.1	Hardware - Input Devices	22
2.2.2	Hardware - Output Devices	26
2.2.2.1	Visual	26
2.2.2.2	Touch	28
2.2.2.3	Auditory	29
2.2.2.4	Smell and Taste	30
2.2.3	Software Tools	30
2.3	The Challenges of Creating an Effective Virtual Environment for Ball Sports	31
2.3.1	Do Sports Skills Reliably Transfer from Virtual Environment Training Conditions to Real-World Scenarios?	32
2.3.2	What Types of Skills Appear to be Best Suited to Training in Virtual Environments?	35
2.3.3	Which Technologies Achieve the Best Results in Specific Performance Measures?	36
2.3.4	Should Stereoscopic Displays be Used?	41
2.3.5	Is High Fidelity Always Better?	44
2.3.6	How, What and When Should Feedback be Delivered to the Learner?	48
2.3.7	When is a Virtual Environment Not Successful in Teaching Motor Skills?	50
2.3.8	What Does it Cost?	52

2.4	Summary	53
3	VE Testbed	55
3.1	System Design	56
3.1.1	Hardware	58
3.2	Creating the Virtual Environment for Rugby Skills Training (VERST)	63
3.2.1	The Ball: tracking and modelling solutions discussed	63
3.2.2	Creating the Ball Model in XVR	63
3.2.3	The Virtual Reality Peripheral Network	64
3.2.4	The Method for Ball Capture	65
3.2.5	Vicon Motion Capture	69
3.3	Constructing the Virtual Rugby Pitch	72
3.3.1	The Pitch	72
3.3.2	The Skybox	74
3.3.3	The Targets	74
3.3.4	The Camera	76
3.3.5	The Flag and Windsock	78
3.4	Additions to the Environment	82
3.4.1	Interactivity	82
3.4.1.1	Hitting the Virtual Targets	82
3.4.1.2	Moving Targets	83
3.4.1.3	Ball Trails	84
3.4.2	Administrative Controls	86
3.4.2.1	System Reset	86
3.4.2.2	Pre-prepared Throw (Debug Mode)	86
3.4.3	User Interface	87
3.4.3.1	Keyboard and Mouse Control	87
3.4.3.2	Statistics Console	88

3.4.3.3	Display Options Console	89
3.5	Rugby Ball Physics Engine	89
3.5.1	The Straight-Line Prototype	90
3.5.2	Velocity Calculations	91
3.5.3	The First Basic Curve	92
3.5.4	Adding The <i>C Set</i>	94
3.5.5	Rugby Ball Dynamics	95
3.5.6	Model Development	98
3.5.7	Adapting the Model for Weather Data	101
3.5.8	Testing the Model using Vicon Data	101
3.6	Summary	104
4	Experiment 1: An Initial Evaluation of the Efficacy of VERST	106
4.1	Description of CRVM Facilities	107
4.2	Experimental Design	109
4.3	Developing the Experiment	112
4.3.1	Porting of the System to Virtools	115
4.3.2	Developing the Configuration File	116
4.3.3	Adaptation of the Physics Engine	116
4.4	The Experiment	118
4.4.1	Experimental Procedure	118
4.4.2	Participants	121
4.4.3	Data Gathered	125
4.5	Results	127
4.5.1	Observations and Ideas Explored	131
4.5.2	Statistical Analysis	133
4.6	Discussion	140
4.7	Summary	141

5	Experiment 2: Further Evaluation of the Depth Perceived in VERST	143
5.1	Description of the Hartree Centre Facilities	144
5.2	Experimental Design	145
5.3	Developing the Experiment	147
5.3.1	Porting of the System to Unity	147
5.3.2	Scripting Interaction	151
5.3.3	Hardware Interaction with MiddleVR	153
5.4	The Experiment	154
5.4.1	Experimental Procedure	154
5.4.2	Participants	158
5.4.3	Data Collected	160
5.5	Results	162
5.5.1	Verbal Task	162
5.5.2	Throwing Task	170
5.5.3	Presence Questionnaire	170
5.6	Discussion	172
5.6.1	Comparing the Results of the Verbal and Throwing Experiments	172
5.6.2	Follow-up Questions for the Participants	173
5.7	Summary	177
6	Discussions and Conclusion	179
6.1	Summary	180
6.2	Comparing the Results of Experiments 1 and 2	182
6.3	Research Questions	184
6.3.1	Do Sports Skills Reliably Transfer from Virtual Environments Training Conditions to Real-World Scenarios?	185
6.3.2	What Types of Skills Appear to be Best Suited to Training in Virtual Environments?	185

6.3.3	Which Technologies Achieve the Best Results in Specific Performance Measures?	186
6.3.4	Should Stereoscopic Displays be Used?	187
6.3.5	Is High Fidelity Always Better?	187
6.3.6	How, What and When Should Feedback be Delivered to the Learner?	188
6.3.7	When is a Virtual Environment Not Successful in Teaching Motor Skills?	189
6.3.8	What Does it Cost?	189
6.3.9	Conclusion	189
6.4	Future Work	190
6.4.1	Further Experiments	190
6.4.2	Improvements to the Perceptual and Functional Fidelity	191
6.4.3	Improvements to the Physics Engine	191
6.4.4	Implementation of Scenarios and Training Activities	191
6.4.5	Attempting to Solve the Poor Results from the Throwing Task	192
	Bibliography	193
	List of Terms	218
7	Appendix	224
7.1	Software Evaluation Tables	225
7.2	IRB Regulations: The Pitch	228
7.3	IRB Regulations: The Ball	230
7.4	Physics Engine Code (XVR)	231
7.4.1	Function: Integrate()	231
7.4.2	Function: Rigid1()	232
7.4.3	Function: Rigid2()	233

7.5	Experiment 1	234
7.5.1	Full Specification	234
7.5.2	Participant Guidelines and Consent Form	236
7.6	Experiment 2	238
7.6.1	Full Specification	238
7.6.2	Configuration Files for Verbal Task	242
7.6.3	Configuration Files for Throwing Task	243
7.6.4	Participant Guidelines and Consent Form	244
7.6.4.1	Throwing Task	244
7.6.4.2	Verbal Task	245
7.6.5	Presence Questionnaire	247
7.6.6	Presence Questionnaire Full Results	251
7.7	Poster Presented at JVRC 2011	252

List of Figures

2.1	Rugby World Cup 2011 game for the Playstation 3 (505 Games, HB Studios).	15
2.2	The forces acting upon an object in motion.	16
2.3	The Magnus Effect acting upon a sphere spinning (a) forward, (b) backwards and (c) sideways.	18
2.4	A histogram showing number of occurrences of different technologies used over 25 published papers discussing ball sports VEs.	22
2.5	Rugby duel simulator (Brault et al., 2009); (a) shows the virtual pitch and (b) the virtual opponent.	27
2.6	Rugby lineout simulator (Brook et al., 2007); (a) and (b) show the user throwing the tracked ball through the TAM device.	36
2.7	Virtual handball; note the blue detection patches in the four corners of the goal; picture courtesy of Bolte et al. (2010).	37
2.8	“Sporting duels” between rugby players; real data is captured to animate virtual opponents. The user wears an HMD and indicates the direction they believe the opponent will move on a gamepad (Brault et al., 2009).	40
2.9	V-Pong table tennis; picture courtesy of Brunnett et al. (2006).	51
2.10	Virtual table tennis for two players; picture courtesy of Li et al. (2010).	52

3.1	(a) The Personal Space Tracker optical tracking device. (b) The retroreflective markers as seen by the infra-red cameras.	59
3.2	(a) A rugby ball covered in retroreflective markers. (b) The model of the rugby ball in PS Track; each point represents a retroreflective marker affixed to the surface of the ball. (c) Three tracked objects in the PS Track software. (d) The three tracked objects in real space. . .	61
3.3	(a) Ball texture—note texture coordinates are reversed so that it maps correctly in the XVR coordinate space. (b) The ball specifications according to law 2 of the IRB game laws.	64
3.4	Skill-tec training ball anchored to the ground with a large flat board. .	66
3.5	(a) Experimental setup of the hardware. (b) Diagram of the experimental setup described in figure 3.5a.	66
3.6	The initial design for the capture system.	67
3.7	New model for data capture, incorporating an extra data set to provide a more accurate representation of the velocity.	69
3.8	Marker-map for the ball during the Vicon capture. The end markers were spherical, side markers 2, 5, 8 and 11 were flat, circular markers, and the remaining markers were flat squares cut from a strip of retroreflective tape.	70
3.9	Matlab plot of the four panel markers and two end markers.	71
3.10	(a) Initial basic texture for the pitch. (b) Posts created with the basic procedural objects available in XVR.	72
3.11	Three images used to create the new pitch texture.	73
3.12	A new, more complex ground texture created from a combination of the textures in figure 3.11.	73
3.13	(a) The free-to-use sky texture used for the skybox. (b) The scene with the newly-added skybox.	75

3.14	(a) Bull's-eye target texture. (b) The source image for the Brian Moore's Mouth target texture. (c) Brian Moore's Mouth target texture.	76
3.15	(a) Plain rugby player target; (b) blanked-out rugby player shaped target; (c) blanked-out rugby player shaped target holding small bull's-eye target; (d) plain rugby player target holding small bull's-eye target.	77
3.16	Each of the different target types present in the scene: (a) the bull's-eye, (b) Brian Moore's Mouth, and (c) the player holding a small target.	78
3.17	The static camera view of the pitch.	79
3.18	(a) The flag and the windsock with no wind in the scene. (b) Illustrates the calculation of the flag direction based on the wind velocity vector.	80
3.19	The target position being used to create a bounding box around the target.	82
3.20	Here, the conditions for intersecting the target bounding box were met on yz -plane (left) but not the xy -plane (right)	83
3.21	The movement schemas in table 3.2, (a) shows the target from the front, (b) from the side.	84
3.22	Ball trails being demonstrated with a tracked throw.	85
3.23	The debug throw shown using ball trails.	86
3.24	(a) The options menu displayed when the right mouse button is clicked on-screen. (b) The statistics console. (c) The settings console.	88
3.25	Planned parabolic flight path of the ball.	93
3.26	Diagrams by Bray and Kerwin (2003).	96
3.27	Matlab implementation of the physics engine in action. The red line represents the data of the real tracked ball; the blue lines each represent a predicted flight path by the physics engine given the real data up to the point from where they start.	102
4.1	The CAVE configuration possibilities.	108

4.2	A top-down view of the screen configuration. The grey sections highlight the screens, whereas the ‘A’ and ‘B’ markers show the two positions the participant stands during the experiment. The ‘V’ markers give the position of the Vicon cameras and the red hatched area indicates the approximate capture area.	109
4.3	(a) The tethered ball suspended from the ceiling with a rope. (b) One of the Vicon cameras used in the experiment.	110
4.4	(a) The tracking rig suspending 5 retroreflective markers at the end of the rugby ball. (b) Infitec glasses with retroreflective markers on the periphery of the lenses. (c) The semi-opened CAVE with the ball suspended from the ceiling of the building and five Vicon cameras surrounding the ball.	114
4.5	The configuration file with example values. Line numbers are shown for reference.	117
4.6	(a) Calibrating the ART system with the calibration wand. (b) The Essilor Digital CRP used to measure participant’s interpupillary distance.	119
4.7	(a) The control station for the experiment, situated next to the CAVE, a technician would be present at the station during the entire experiment. (b) The view of the VERST system from the control station. The CAVE walls are marked by the grey checkered planes. The trial details are displayed in the top-left corner of the window.	120
4.8	An axis of spin down the centre of the ball (a) against an axis of spin at approximately 25° causing a skewed spin (b).	123
4.9	Example output file from experiment 1.	126
4.10	Example trial data from experiment 1.	126
4.11	The failed trials from table 4.3 plotted as a stacked bar chart; this highlights the large number of capture failures compared to timed out trials.	128

4.12	(a) Participant 1, configuration 1, trial 1 visualised in Matlab. The large red circle is the target; the small blue circle is participant 1's head position; the green square shows the location of the floor screen; the purple section of the flight path marks the captured real data while the yellow portion marks the resulting flight path calculated by the physics engine. (b) The $y-z$ plane of (a). (c) The $x-z$ plane of (a).	130
4.13	Each configuration is individually plotted (c1–8). The diagonal dashed line denotes the 1:1 ratio of the two axes. (a) For all participants, the average distance thrown as a function of target distance. (b) Participants 1–3, the average distance thrown as a function of target distance. (c) Participants 1–3, the average distance thrown as a function of target distance from figure 4.13b presented as a series of best-fit lines. . . .	134
4.14	Average distance thrown by participants 1–3 across all configurations as a function of the target distance. Each graph only contains the data from the trials representing the variable condition designated in the title. (a) Configurations 1–4; (b) configurations 5–8. The dashed blue lines denotes the 1:1 ratio of the two axes.	135
4.15	Average distance thrown by participants 1–3 across all configurations as a function of the target distance. Each graph only contains the data from the trials representing the variable condition designated in the title. (a) Configurations 1, 2, 5 and 6; (b) configurations 3, 4, 7 and 8. The dashed blue lines denotes the 1:1 ratio of the two axes.	136
4.16	Average distance thrown by participants 1–3 across all configurations as a function of the target distance. Each graph only contains the data from the trials representing the variable condition designated in the title. (a) Configurations 1, 3, 5 and 7; (b) configurations 2, 4, 6 and 8. The dashed blue lines denotes the 1:1 ratio of the two axes.	137

5.1	(a) A virtual car cockpit, featuring a real car seat and steering wheel with a virtual exterior scene. (b) Testing the curved screen using the testbed environment.	145
5.2	(a) and (b) The original stadium model. (c) The grass image used to texture the ground of the pitch. (d) The full pitch environment.	149
5.3	(a) The original goal post model; the posts are sized based on the IRB regulations. (b) Re-sized posts to counteract the z-fighting. (c) The textured target model in Unity. (d) The textured ball model in Unity.	150
5.4	The scene from the viewpoint of the user.	150
5.5	The instruction text and countdown displayed in the centre of the screen.	152
5.6	Example output file of the verbal task in experiment 2.	152
5.7	(a), (b), and (c) are random dot stereograms to be viewed with anaglyphic glasses (d).	155
5.8	The IPD ruler: to use, hold to subject's face and measure the centre for each pupil	156
5.9	Participants during the experiment performing the (a) throwing and (b) verbal tasks.	157
5.10	Results of the verbal task. Intended perceived distance between 5–10 m is plotted against the distance estimated by all participants. The dashed line represents the 1:1 ratio of the two axes.	163
5.11	The results of the verbal task in experiment 2 with the modified results of participant 4. The dashed line representing the 1:1 ratio of the intended to estimated distances.	164
5.12	Results of the verbal task separated into the estimations made under different configurations (refer to table 5.1 and table 5.2 for configuration descriptions) using the curved and flat screens. The horizontal dashed lines mark the range of target distances intended to be perceived (5–10 m).	165

5.13	Average distance verbally estimated by participants 1–6 across all configurations as a function of the intended target distance. Each graph only contains the data from the trials representing the variable condition designated in the title: (a) configurations 1–3; (b) configurations 5–6. The dashed pink lines denotes the 1:1 ratio of the two axes. . . .	166
5.14	Average distance verbally estimated by participants 1–6 across all configurations as a function of the intended target distance. Each graph only contains the data from the trials representing the variable condition designated in the title: (a) Configurations 1 and 4; (b) configurations 2 and 5. The dashed pink lines denotes the 1:1 ratio of the two axes.	167
5.15	Average distance verbally estimated by participants 1–6 across all configurations as a function of the intended target distance. Each graph only contains the data from the trials representing the variable condition designated in the title: Configurations 3 and 6. The dashed pink lines denotes the 1:1 ratio of the two axes.	168
5.16	The estimated distance minus the physical distance of the participant from the screen is plotted against the intended perceived distance. The dashed line representing the 1:1 ratio of the intended to estimated distances.	169
5.17	Distance thrown according to the results of the physics engine plotted against the intended target distance. The dashed line represents the 1:1 ratio of the two axes.	171
5.18	Estimated distance during the verbal (dashed line) and the throwing (solid line) experiments plotted against the intended distance. The lighter-coloured dashed line represents the 1:1 ratio of the two axes. .	173

6.1	Comparison of the average estimated distance between the verbal and throwing tasks in the two experiments as a function of the intended distance. The lighter-coloured dashed line represents the 1:1 ratio of the two axes.	183
7.1	Ground specifications according to law 1 of the IRB game laws	228
7.2	Ball specifications according to law 2 of the IRB game laws.	230
7.3	Function to approximate the next position of the ball using RK4.	231
7.4	Function to evaluate $k1$ using equations 3.22, 3.23 and 3.24.	232
7.5	Function to evaluate $k2$, $k3$ and $k4$ using equations 3.22, 3.23 and 3.24.233	
7.6	Trial timings are read from this file.	242
7.7	Participant number is used to number the output files, while the configuration order is read into the program and interpreted to set order the configurations are called.	242
7.8	Trial timings are read from this file, along with tracker settings.	243
7.9	Participant number is used to number the output files, while the configuration order is read into the program and interpreted to set order the configurations are called.	243

List of Tables

3.1	The average data readings for the x coordinate values of segments of the capture area (the left of capture area is physically to the right of the camera which is facing the user, and vice versa for the right of the capture area).	68
3.2	The movement schemas available for targets.	84
3.3	Key-presses versus the function or task that they trigger.	87
4.1	Trial Configuration table for experiment 1.	113
4.2	Details of the participants recruited for the experiment.	122
4.3	Number of failed trials during the experiment which required a reattempt.	128
4.4	Results of correlation testing between variables 1 and 2 (sorted alphabetically by variable 1).	132
5.1	Trial configuration table for the verbal task in experiment 2.	146
5.2	Trial configuration table for the throwing task in experiment 2.	146
5.3	The MiddleVR settings for the curved and flat screens in experiment 2.	153
5.4	Details of the participants recruited for experiment 2.	158
5.5	Descriptive statistics of the verbal experiment, plotted in figure 5.10.	163
5.6	Median verbal estimations by participant position, minus the physical distance between the participant and the screen when they are in that position.	168

5.7	Descriptive statistics of the verbal experiment: comparing the original data to the data minus the participant's physical distance from the screen.	169
5.8	Presence Questionnaire Group Results	170
5.9	Participant's responses to the four follow-up questions asked to them.	174
5.10	Participant's responses to the fifth follow-up questions asked to them.	176
7.1	Software for interfacing between the hardware and the VE.	225
7.2	Software for creating the virtual scene.	226
7.3	Full Configuration Procedure of VERST for Experiment 1	234
7.4	Full Configuration Procedure of VERST for Experiment 2	238

Acknowledgements

I would particularly like to thank my supervisors Professor Nigel W. John and Dr Serban Pop for guiding me through the project and making sure I kept on target, and to my steering group Dr Simon Watt, Dr Gavin Lawrence and Kenton Morgan for helping me to understand the psychological and sports science sides of the research better.

Thank you to my parents who encouraged me to keep going, and to Thom for putting up with me when I got stressed.

Many thanks to research collaborators Daniel Mestre, Vincent Perrot and Pierre Mallet for their hospitality and support during the experiment at CRVM. Thanks also to Terry Hewitt from the Hartree Centre, and Joss Petit and Chris Greive from OPTIS Northern Europe who supported the experiments at the Science and Technology Facilities Council, Daresbury, UK.

This research was funded by a Knowledge Economy Skills Scholarship (KESS), with contributions from the European Social Fund (ESF), Bangor University and Rygbi Innovations Cyf. The research undertaken at Centre for Virtual Reality of the Mediterranean (CRVM) at Aix-Marseille University, France was partially supported by the European Community's Research Infrastructure Actiongrant agreement VISIONAIR 262044—under the 7th Framework Programme (FP7/20072013).

Chapter 1

Introduction

1.1 Context and Motivation

Professional sport has become so competitive that trainers and coaches are seeking new methods of improving their elite athletes that will give them a step up above their competitors. In recent years, the advancement of technology has reached a level of sophistication that has brought it to the attention of those seeking additional methods of training. The use of motion capture systems allows coaches and trainers to observe the movement of their athletes in great detail, and at low enough speeds to detect any flaws in their actions.

The company Rygbi Innovations Cyf. are interested in exploring the potential of virtual environments to assist in the training of rugby union players. They are particularly interested in a system which could be implemented at low cost with easy-to-use hardware and software, with the aim of introducing the technology at a high level (larger teams first) before disseminating to smaller clubs that will not have the budget or technical expertise available to the larger teams. Rygbi Innovations Cyf. were not focussed on a particular skill to be trained, as they were interested in what could be accomplished overall with the available technology.

This research examines a candidate system developed for training players in the skill of passing using a virtual environment.

1.2 Hypothesis

The hypothesis investigated in this thesis is: “A virtual environment can be used to effectively train ball passing skills in rugby.”

We believe that it is possible to develop a virtual environment that will positively assist in training ball passing skills to rugby players. We also believe that such an environment can be adapted for training and analysis of a variety of other sports skills.

1.3 Contributions

This thesis explored the use of a virtual environment to act as a training resource for training passing skills in the sport of rugby union. A testbed environment exploring a system in which the user must throw a real ball towards a virtual target has been designed, and, to our knowledge, is the first to do so. The experiments conducted in this research suggest that the users are unwilling to throw a real ball towards expensive equipment at full force; therefore, the task of throwing the real ball may not be a suitable method of engaging a user in training throwing skills.

1.3.1 List of Publications

Sections of the work presented in this thesis have been published.

The main author of these publications was Helen Miles. S. R. Pop, S. J. Watt, G. P. Lawrence and N. W. John contributed with their individual expertise to plan and develop the two experiments and analyse the results. V. Perrot, P. Mallet and D. R. Mestre, from the collaborating team at Aix-Marseille University, provided the facilities and their expertise for the development and running of the first experiment.

Poster Publications

H. C. Miles, N. Musembi, S. R. Pop, & N. W. John, "A Virtual Environment for Rugby Skills Training," In *Proceedings of the Joint VR Conference of EuroVR and EGVE (JVRC 2011)*, Nottingham, UK. VTT, Vuorimiehentie, Finland, 2011, pp. 64–66.

The poster can be seen in appendix 7.7.

Journal Publications

H. C. Miles, S. R. Pop, S. J. Watt, G. P. Lawrence & N. W. John, "A review of virtual environments for training in ball sports," *Computers & Graphics*, 36(6), 2011, pp. 714–726.

This publication contains section contributions from each of the other authors, based on their individual expertise. A large portion of this publication has been used in chapter 2 of this thesis.

H. C. Miles, S. R. Pop, S. J. Watt, G. P. Lawrence, N. W. John, V. Perrot, P. Mallet, D. R. Mestre & K. L. Morgan, "Efficacy of a Virtual Environment for Rugby Skills Training," *LNCS Transactions on Computational Science Special Issue on Cyberworlds*, 2014, In press.

Portions of this publication have been used in chapter 4 and chapter 5 of this thesis.

Peer Reviewed Publications

H. C. Miles, S. R. Pop, S. J. Watt, G. P. Lawrence, N. W. John, V. Perrot, P. Mallet & D. R. Mestre, "Investigation of a Virtual Environment for Rugby Skills Training," In *Proceedings of the 2013 International Conference on Cyberworlds (CW 2013)*, Yokohama, Japan. pp. 56–63.

A portion of this publication has been used in chapter 4 of this thesis.

1.4 Thesis Structure

Chapter 2 will explore the first concepts of the project, including a brief history of virtual environments for general purposes, an examination of the software and hardware used to create these environments and concluding with an section detailing some of the projects dealing specifically with virtual environments for ball sports. From this literature, a number of questions surrounding the project are raised and carefully considered.

Following a review of the current literature, the design and development of the testbed environment is explored in detail in chapter 3. There will be a focus on each area of the system: the virtual scene, the tracking technology, the physics model managing the ball movement and the user interface and human-computer interaction aspects.

The efficacy of the system was tested in two experiments. Chapter 4 details an experiment to evaluate the initial performance of VERST, which was conducted in a CAVE-like environment at Aix-Marseille Université (Marseille, France). Chapter 5 describes an experiment to further evaluate the depth perceived in VERST, conducted using two power walls at the Science and Technology Facilities Council (Daresbury, UK). The results of each experiment are discussed individually at the end of their respective chapters, and the experiments are compared at the end of chapter 5.

In chapter 6, the VERST system is discussed in the context of the research questions asked in chapter 2. Some potential future work to be done on VERST is considered based on the issues found during chapter 4 and chapter 5.

Chapter 2

Background

This chapter aims to illustrate the current state-of-the-art technologies and research problems involved in improving functional fidelity for the trainee, where functional fidelity pertains to the similarity of the athlete's behaviour in the simulated and real environments. A background to virtual environments follows, including a brief overview of the contribution of sports-themed video games (many innovations stem from this sector).

The forces experienced by sports projectiles are described, followed by the explanation of an example system of motion derived to describe the motion of a rugby ball in flight.

Clearly, the utility of a virtual environment for sports training is dependent on the capabilities of the available technology components. Thus, a summary of the hardware and software technologies required to build a high functional fidelity sports-themed virtual environment is then presented. These are explored in more detail in the context of the challenges that must be overcome to build an effective training environment for ball sports, and the advantages that can be achieved are highlighted. Examples from current multi-disciplinary research are used to demonstrate each point. The chapter concludes with suggestions for future research directions.

2.1 Background: First Concepts

In this first section the roots of sports-based virtual environments will be examined in depth: an essential introduction to the field of motor skill learning is the starting point, followed by an introduction to general and sports-based virtual environments, a discussion of the relevance and influence of sport-themed computer games and finally a survey of the more prevalent technology components in sports virtual environments.

2.1.1 Introduction

There is growing interest in utilising Virtual Environments (VEs) in the context of sports. In particular there is a desire to be able to improve sensorimotor skills rather than just using a VE as a tool for strategy analysis, or entertainment. The range of skills required across all different sports is very large and varied. A VE needs to provide realistic rendering of the sports scene to achieve good perceptual fidelity. More important for a sport-themed VE is high functional fidelity, which requires an accurate physics model of a complex environment, real time response, and a natural user interface. The goal is to provide multiple scenarios to players at different levels of difficulty, providing them with improved skills that can be applied directly to the real sports arena.

2.1.2 Motor Skill Learning

The way people learn and execute effective motor skills poses a diverse and challenging set of questions that have been the attention of researchers since the 19th Century (Woodworth, 1899). However, despite over 100 years of research on skill acquisition, the fundamentals of effective coaching have still not been fully illuminated. What is clear is that in order to develop expert performance, individuals must invest considerable periods of practice time (e.g., a 10 year period, as suggested by Simon and Chase (1973), and Starkes et al. (1996)), and that this practice should be specific to the skills

required for the to-be-learned task/sport (i.e., sports-specific or deliberate practice, see Ericsson et al. (1993); for a review see Starks (2000)). The body of research has identified that both coaching and skill acquisition normally involve three distinct processes: conveying information (i.e., observational learning), structuring practice (i.e., contextual interference) and the nature and administration of feedback (i.e., feedback frequency, timing and precision) (Lawrence and Kingston, 2008). Additionally, the time constraints placed on both coaches and performers require the teaching and learning of more than one skill in a single session. Consequently, most training sessions demand that learners practise a wide range and number of skills. Indeed the skill acquisition literature suggests that learning is typically maximised when performers are forced to learn a variety of skills in a single training session and where these skills are practised in a random fashion (i.e., the practice environment is very diverse and challenging, similar to the demands of a competition or game) (Battig, 1966, 1979; Brady, 2004; Lee and Magill, 1985; Perez et al., 2005; Vickers, 2007). This complexity and variety of training is difficult to achieve in conventional training sessions. Moreover, it can be very inefficient because the participation of many players is often required simply to simulate the competitive context being taught, even though only one player might benefit directly from the training.

2.1.3 Virtual Environments

A VE can be considered as a collection of technologies that allow people to interact efficiently with 3D computer generated models in real time using their natural senses and skills (Burdea and Coiffet (2003); Slater et al. (2002)). As such, a VE can be used in a variety of ways, including: to provide the end user with an immersive experience for entertainment purposes; to design a new concept or product; to explore and gain insight from data in an intuitive manner; and to train people at complex (and perhaps dangerous) tasks in a safe, convenient, and/or effective environment. Early adopters of VEs include the aeronautics industry (flight simulators), the entertainment industry

(simulator rides, video games, and movies), medical-procedures training, oil and gas exploration, and automotive design.

As the available hardware gets more powerful, and often less expensive, more and more application areas are taking advantage of VEs. Sport performance is one area where there has recently been great interest in obtaining added value from a VE. The effectiveness of any VE used for training can reasonably be measured against three criteria: it must be realistic enough, it must be affordable, and it must be validated to show that it works as designed. It is possible that it is only by satisfying all of these criteria that a VE built for training skills in sport will become an accepted tool.

2.1.4 Virtual Environments for Sports Training

There are many potential advantages to a sports training VE. For example, a VE allows standardised scenarios to be created by coaches; extra information can be given to players to enrich the environment and guide their performance; and it is possible to quickly and randomly change the environment to match any competitive situation (including simulating the presence of other players). All of these have the potential to revolutionise current training practices. The vast range of skills, environments, and levels of training encountered across all sports is beyond the scope of a single PhD project. The focus here is on ball sports, because this area is most relevant to the current goal of producing a training VE for aspects of rugby. Specifically, there is a focus on motor-skill learning, but it should be noted that VEs also have the potential to provide a sophisticated tool for training in strategy and tactics (Hughes and Franks, 2004).

2.1.5 Validation and Verification

Verification and validation of a system are ways in which a system must be examined to determine if the performance is adequate. Validation is the analysis of a simulation

system to examine whether it is an accurate depiction of the real system it is attempting to simulate. Verification is the analysis of the functionality of a system to determine whether it works as it was intended to (Law et al., 1991).

Validation and verification (V & V) are an important consideration in the creation of a VE in a number of ways:

- validity of the model being simulated;
- whether skills transfer can occur from the VE to the real world; and
- basic software validity.

For a VE which simulates a real-life system, for example a physical or biological system, it is imperative the system be proven to be accurately portraying the real-life system before it can be relied upon to provide correct information (Kleijnen, 1995). Balci (1997) states that to test the accuracy of a simulation, several assessment activities should be undertaken. The model should be verified to determine whether it has been converted from the original derived formulae or flowchart into a program accurately. The model should be validated to determine if it “behaves with satisfactory accuracy” with reference to the objectives of the simulation – is it the correct model for the purpose? The model should then be tested with real data to search for any errors or inaccuracies existing within it.

Balci (1995) offers fifteen principles for performing V & V on a simulation, including that the outcome of V & V testing should not be considered simply as ‘success’ or ‘failure’, but as “a degree of credibility”. That the “model credibility can be claimed only for the prescribed conditions for which the model is tested”, is it not possible to completely (exhaustively) test the model and “successfully testing each submodel does not imply overall model credibility”. He also provides a taxonomy (Balci, 1994) of techniques that can be utilised during the development of the simulation model, and states that the model should be continuously undergoing evaluation using these techniques.

Sargent (1998) gives details of a large number of techniques for V & V purposes, among which are:

- comparison to other models;
- event validity – whether events in the simulation comparable to the reality they are simulating;
- face validity – asking experts on the reality being simulated whether the behaviour of the simulation is “reasonable”;
- use of historical data, and whether the simulation produces the same data as the reality it is modelling; and
- internal validity – a determination of how variable the results produced by the system can be, or how consistent it is.

For VEs that attempt to train a skill or task, operational validity – whether “the model’s output behaviour has the accuracy required for the model’s intended purpose” (Sargent, 1998) – and skills transfer is extremely important. The VE must be proven to effectively enhance the user’s ability to perform a task before it can be used as a training tool. Given that psychological experiments are often used to evaluate the performance of users, the concept of validity in psychological experiments should also be noted. In psychometric terms – that is, the study of measurement in psychological experiments – “validity refers to the degree to which evidence and theory support the interpretations of test scores entailed by proposed uses of tests” (American Educational Research Association et al., 1999). Messick (1995) defined six aspects in psychometric validation:

1. content validation, where it is ensured that the form of testing is relevant to what the experimenter wishes to evaluate;

2. substantive validation – that the participant is able to perform the task to be measured;
3. structural validation, where the structure with which the experiment is scored is evaluated to ensure it measures the relevant results;
4. generalizability, whether the experiment can be generalized to examine larger groups and different tasks?;
5. external validation, whether the experiment has convergent, discriminant, and predictive qualities; and
6. consequential validation, where the risk of inappropriately interpreting the results is evaluated.

As an illustration, Schwebel et al. (2008) attempted to validate their VE to examine child pedestrian injury while crossing roads. Recruiting 74 adults and 102 children to take part in their experiment, they examined the behaviour of the participants in a real environment, and a simulated version of the real environment. They examined the VE in a number of ways:

1. Use of a custom questionnaire to determine how real the participants deemed the environment to be;
2. measured statistics of the task performance in the VE on a number of variables (number of hits, close calls, average temporal gap between the time the participant safely crossed the street and the arrival of the next vehicle into the crosswalk, wait time by traffic density, and start delay);
3. performed correlation testing on the variables to ensure internal reliability of the results;
4. testing the virtual responses against real world behavioural data to ensure the participants reacted to the task in a similar manner in both; and

5. parent-reported measures of child temperament to examine the child's regular behaviour and determine whether their responses were related to behavioural traits.

Based on their findings, Schwebel et al. believed their VE to be a valid method for examining child pedestrian injury.

A common form of testing whether the user has successfully learnt from a VE is to evaluate the performance of a group of users both in reality and an equivalent virtuality, and to compare those results. This method has been used to evaluate several VEs that aim to train laparoscopic surgery (also known as minimally-invasive, or keyhole surgery) (Grantcharov et al., 2004; Arora et al., 2010; Seymour et al., 2002). Each has demonstrated that the participants in the experimental group using the VE for training were more confident and performed the surgery faster than those without VE training. Similar methods can be employed to examine the validity of skills transfer in sports VEs, as it is possible to measure the performance of a sportsperson in a specific task in reality and an equivalent virtuality. This will be discussed with example systems later in the chapter.

For any software, V & V of the development process should also be considered. For software purposes, the terms can be thought of as "Validation: Are we building the right product?" and "Verification: Are we building the product right?" (Boehm, 1979). An excellent example of V & V of the software behind a VE can be seen in Pioch's report on the naval training VE 'Officer of the Deck' (Pioch, 1995). Pioch notes that there are usually five processes for verifying and validating simulation software (Knepell and Arangno, 1993):

1. conceptual model validation,
2. software verification,
3. operational validation,

4. data validation, and
5. internal security validation.

His evaluation of the Officer of the Deck system was concentrated on items 2–4 of this list. His definitions for conceptual model validation and software verification match those previously defined, whilst also discussing three further concepts. Operational validation, in which the software is tested to ensure that “the model compares well to perceived reality”. Data validation, where the “source and consistency of the data used in the computerized models” is checked. And finally internal security validation, where the system is tested to determine if it is sufficiently protected from the threat of viruses and malicious access.

2.1.6 Sports-themed Computer Games

At first thought, computer gaming might seem unlikely to inform VEs for professional sports training. Such games can be considered as a low functional fidelity VE, however, and although designed primarily for entertainment, many useful tools and techniques have been developed in this domain that can be deployed in training applications. These range from hardware peripherals, to physics engines, to computer-graphics rendering algorithms.

For example, a significant evolution within the computer games market followed the introduction in 2006 of the Nintendo (Kyoto, Japan) Wii, which uses a wireless remote interaction device—the Wii Remote—with built-in accelerometers and infrared sensors to detect three-dimensional movements. It allows users to play and control games using physical gestures combined with traditional button presses. Accuracy is increased with the Wii MotionPlus (Nintendo, 2014) add-on, which captures rotational motions using two gyroscopes and an angular rate sensor, helping to disambiguate complex movements. The Wii Balance Board is also available, which can measure the weight of the player and his or her centre of gravity, and is specifically aimed

at fitness activities. Game players can therefore now improve their general fitness as they engage in the game activity. The popular Wii Sports package provides simplified versions of tennis, baseball, bowling, golf, and boxing, and evidence is emerging that the Wii platform is contributing to the acquisition of skills in some real sports (e.g. 10-pin bowling, Siemon et al. (2009)).

Another recent innovation in games hardware is the Microsoft (Redmond, USA) Kinect for Xbox 360 and Windows, which employs camera-based marker-less tracking to encode the user's movements. The images from an RGB camera inside the Kinect's sensor bar are used to detect moving people against background objects and, together with the integrated depth sensor, can support human-feature recognition. The Kinect can build an accurate motion analysis map, and can track up to 20 separate points of articulation for each person in its range. The potential of this hardware is demonstrated by the Kinect Sports package that includes beach volleyball, boxing, bowling, soccer, table tennis and athletics. The user remains in the same place but mimics the actions and gestures of a real athlete using arm and leg movements. The Kinect may require calibration at the start of a session, but the accuracy achieved in play is good (considered better than on the Wii). However, independent reviews also report problems with the responsiveness of the device and general game play (Kotaku, 2012).

Video games continue to provide innovative technology at an affordable price. They can also excel at recreating the actual environment of a sport, from a high-quality reproduction of the stadium to the realistic scale of all the elements in play (balls, bats, gymnastic apparatus, etc.). Some games are designed as a first-person experience, to enhance the feeling of immersion, but others effectively re-enact a television broadcast by featuring realistic play-to-play audio commentary and television-specific camera-angles, giving the player a third-person experience (for example, as in the Rugby World Cup 2011 game, see figure 2.1; Push Square (2014)). The drawback of this latter approach is that often the event flow is suspended whilst the user makes several choices

about how he or she wants to continue. This would not be appropriate in a VE designed to provide skills training that would transfer to the real environment.



Figure 2.1: Rugby World Cup 2011 game for the Playstation 3 (505 Games, HB Studios).

2.1.7 Modeling the Physical Forces

2.1.7.1 The Principles of Motion

The motion of an object must follow Newton's Laws of Motion (Weisstein, 2014b), which are defined as follows:

1. (Law of inertia): A body at rest remains at rest and a body in motion continues to move at a constant velocity unless acted upon by an external force.
2. A force F acting on a body gives it an acceleration a which is in the direction of the force and has magnitude inversely proportional to the mass m of the body:

$$F = ma.$$

3. Whenever a body exerts a force on another body, the latter exerts a force of equal magnitude and opposite direction on the former. This is known as the weak law of action and reaction.

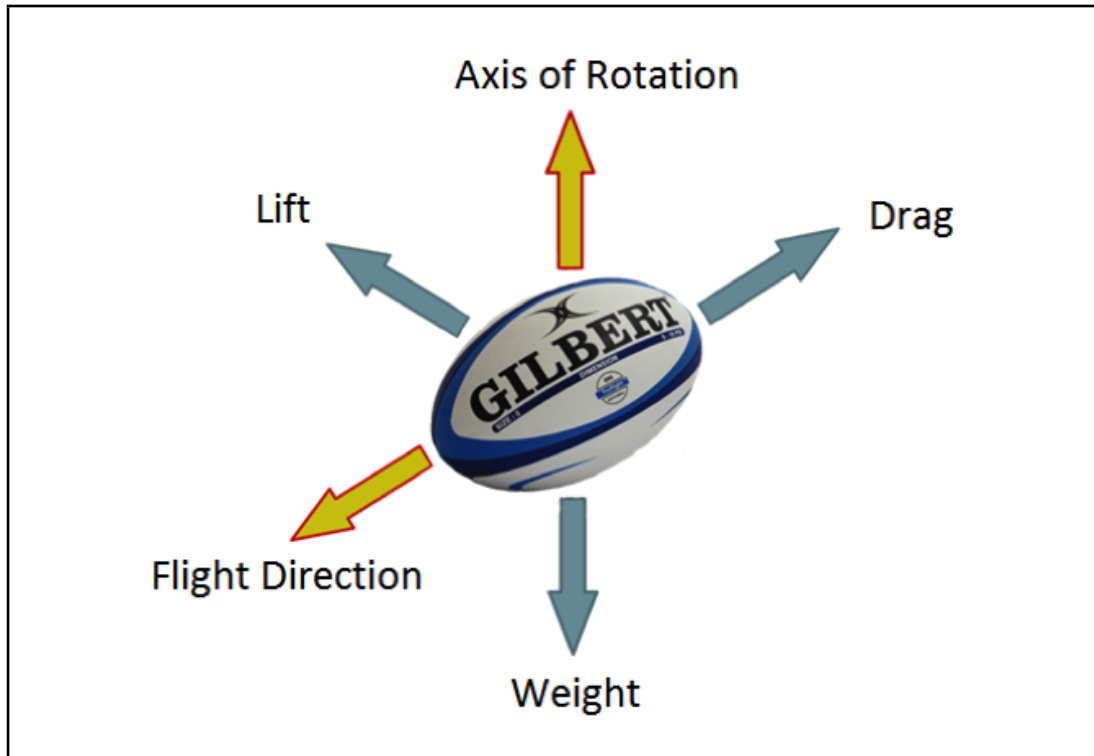


Figure 2.2: The forces acting upon an object in motion.

There are a number of forces acting upon a projectile (an object which received only an initial force to propel it into the air, and is subsequently only acted upon by external forces), as shown in figure 2.2. The projectile moves in a direction according to an initial force; given no other external forces, the projectile would move continuously in the direction of the throw (according to Newton's 1st Law). External forces which affect the flight path of the projectile are the weight of the object in motion, and the drag and lift of the air on the object; the rotation of the object also affects the flight path.

The weight (NASA Glenn Research Center, 2014d) (w) of the projectile is defined as the mass (m) of the object multiplied by gravity ($g = 9.8 \text{ ms}^{-2}$), as shown in equation 2.1.

$$w = mg. \quad (2.1)$$

The drag (NASA Glenn Research Center, 2014b) force (F_D) acting on the projectile is dependent on the surface area (A) of the object that meets the air (density ρ) through which it is travelling, as well as the velocity (magnitude v) at which the object is travelling and a drag coefficient (C_D):

$$F_D = \frac{C_D \rho v^2 A}{2}. \quad (2.2)$$

The drag coefficient is often determined experimentally using a wind tunnel, but can be calculated by rearranging equation 2.2:

$$C_D = \frac{2F_D}{\rho A v^2}. \quad (2.3)$$

The lift (NASA Glenn Research Center, 2014c) force (F_L) on the projectile, like the drag force, is calculated from the air density (ρ), cross-sectional area (A), velocity of the projectile (magnitude v) and the lift coefficient (C_L), which is also determined experimentally using a wind tunnel:

$$F_L = \frac{C_L \rho v^2 A}{2}; \quad (2.4)$$

$$C_L = \frac{2L}{\rho A v^2}. \quad (2.5)$$

If the projectile is spinning, and the air flowing along a curved portion of its body (if it is spherical or cylindrical), the direction of the spin against the air can cause a change in the direction of the projectile; this is known as the Magnus Effect. If the projectile is spinning forwards perpendicular the direction it is travelling (see figure 2.3a), the projectile will curve towards the ground and fall faster, while spinning backwards (see

figure 2.3b) will increase the lift and cause the projectile to travel farther; side spin on the projectile will also cause it to swerve sideways (see figure 2.3c).




Figure 2.3: The Magnus Effect acting upon a sphere spinning (a) forward, (b) backwards and (c) sideways.

As the projectile moves through the air surrounding it, the air is disturbed and turbulent forces occur. The density and viscosity of the air affect the movement of the projectile by producing higher drag when the air is denser and more viscous, and less drag when the air is thinner and less viscous. The ratio of the viscous forces of the air against the inertial forces of the projectile is represented by the dimensionless Reynolds Number (NASA Glenn Research Center, 2014a). The Reynolds Number (Re) is defined as the velocity of the projectile (v) multiplied by the air density (ρ) and the length travelled through the air (L), all divided by the viscosity of the air (μ):

$$Re = \frac{v\rho L}{\mu}. \quad (2.6)$$

2.1.7.2 Physics Engines

A physics engine is the attempted replication of a real-life physical system inside a computer simulation. High-performance scientific simulations, such as Computational Fluid Dynamics (CFD) and Molecular Dynamics (MD) simulations require large amounts of processing power, and can require supercomputers (for example, Anton, Shaw et al. (2008)) to solve.

Rigid and soft body dynamics are two types of physics engine commonly used for computer graphics, and the film and game industries. Rigid body dynamics simplifies the problem of complex movement and collisions between objects by assuming the objects do not deform on contact with each other or when under the influence of forces; this allows simulations to run in real-time, which is ideal for games. Conversely, soft body dynamics allows the deformation of the objects to occur, producing a visually-accurate simulation; soft body dynamics are often used for hair and cloth simulations.

The physics engines used in sports games are becoming very sophisticated, with game production companies investing a great deal of time into the development of complex physics modelling. For example, a team at EA Sports (Electronic Arts Inc., Redwood City, USA) recently spent a year exploring the projectile motion physics used in the FIFA series that was causing unrealistic ball flight motion, discovering that an incorrect drag coefficient was being used (Chiaet, 2013). EA Sports' Player Impact Engine (2014a), developed for their football game FIFA 12, was developed over 2 years to process complex inter-player collisions. More recently, the Player Impact Engine has been adapted into Collision Physics (EA Sports, 2014b), an engine for processing the inter-player collisions in NHL® 14 incorporating the speed and momentum of the players.

Some notable non-sports games also include complex physics simulations, for example Kerbal Space Program (Squad, 2014) features a fully-developed Newtonian dynamics simulation to represent space flight (Groen, 2013), From Dust (Ubisoft, 2014) contains real-time simulations of natural environments, and Universe Sandbox (Giant Army, 2014) simulates the effects of gravity on celestial objects.

2.1.7.3 Modeling the Motion of Sports Projectiles

Sports which have been studied include tennis (Cooke, 2000; Mehta and Pallis, 2001; Cross, 1987), cricket (Mehta and Koga, 2000), golf (Burglund, 2011; Mehta, 1985), football (Alam et al., 2010; Spathopoulos, 2011; Craig et al., 2006; Goff and Carré,

2010; Barber et al., 2009; Griffiths et al., 2005; Barber et al., 2006), baseball (Mehta, 1985; Alaways, 1998), American football (Alam et al., 2011; Rae, 2003), Australian Rules football (Alam et al., 2003, 2009), and rugby (Djamovski et al., 2012; Seo, 2004; Vance et al., 2012; Alam et al., 2008; Fuss and Troynikov, 2012; Takahashi and Kobayashi, 2006; Seo et al., 2004; Djamovski and Alam, 2011; Seo et al., 2006, 2010) (sometimes considered together because of similar shape).

2.1.8 Components Required for a Sports Virtual Environment

There are a number of basic technological elements required to construct a sports VE.

The main technologies typically deployed are:

- Display technologies, often using stereoscopy. Examples range from Head Mounted Displays (HMDs) to large-screen Power Walls;
- Tracking technologies, such as optical, magnetic, and inertial systems. Motion capture is a particular example;
- Haptic technologies for force and tactile feedback. These range from hand-held joystick devices, to large-scale pneumatic motion platforms;
- Software algorithms for efficient rendering of 3D computer graphics and multi-user synchronisation;
- Many custom built devices will also be fabricated depending on the application area.

These components are mainly used in academic projects although there are also a few sports-related commercial systems. Visual Sports (Ontario, Canada), for example, offer relatively advanced simulators that cover golf, soccer, baseball, and other sports. Golf simulators are also available from VirtuGolf (Alabama, USA) and Indoor Golf & Racing (Washington, USA). Their main market is the entertainment industry

but the difference from a video game is that the product is built to be life size and uses the above display and tracking technologies. Sports training is of interest to the commercial sports simulator manufacturer, but currently the majority of VEs available for sports training have been developed by academic groups. Figure 2.4 summarises the hardware components that are typically used as identified from a survey of 25 academic papers (Bideau et al., 2003, 2010a, 2004a,b; Bolte et al., 2010; Brault et al., 2009; Brook et al., 2007; Brunnett et al., 2006; Chong and Croft, 2009; Chong et al., 2008; Chua et al., 2002, 2003; Li et al., 2010; Matsuura et al., 2006; Rusdorf et al., 2007; Syamsuddin and Kwon, 2010, 2011; Vignais et al., 2009a,b, 2010b,a; Xu et al., 2009; Bideau et al., 2010b; Brault et al., 2010; Chung et al., 2010; Gourishankar et al., 2007; Göbel et al., 2010) relating to sports VEs. Many different technologies are used, some clearly favoured more than others, and they are described in the next section of the chapter. The software options most used are also summarised. Key issues and challenges in delivering a ball sports VE are then identified and examples from published work are used to illustrate each point. Finally, possible research directions in the field are discussed.

2.2 Hardware and Software Used for Sports Virtual Environments

The typical hardware and software components needed are identified in the section, and important psychological factors that should be considered are discussed. The challenges that need to be overcome are identified and illustrated with examples from recent projects in this area. The section is divided into input and output devices for hardware, and a discussion of common software packages used.




Figure 2.4: A histogram showing number of occurrences of different technologies used over 25 published papers discussing ball sports VEs.

2.2.1 Hardware - Input Devices

Suitable input devices are required to allow the user to interact and engage with the virtual world in a natural manner. Typically that means that input data need to be processed in real time. Joysticks, keyboards and game pads have all been commonly used in this capacity, for example, by Brault et al. (2009, 2010). For many sports-training applications, and motor-skill learning in particular, however, the VE would ideally encode more complex aspects of movements.

One approach is to use motion capture: the practice of dynamically recording the positions in 3D space of various pre-defined points on the body or an object. There are various mature technologies employed to do this. Optical tracking devices use multiple infrared-sensitive cameras to track the positions of ‘markers’ attached to the target (to joints on the body, for example). ‘Passive’ systems illuminate the scene with infrared light and detect the positions of retroreflective markers. ‘Active’ systems use infrared

light-emitting diodes as markers, but are otherwise similar. Provided the positions of the cameras in the space are known (i.e. the system is calibrated), the 3D location of each marker can be recovered by analysing its position in each camera's image (at least two are required to recover the 3D position of a marker). A quite diverse array of systems is available from different optical motion-capture companies, including Vicon (Los Angeles, USA), OptiTrack (NaturalPoint, Corvallis, USA), PhaseSpace (San Leandro, USA), Qualisys AB (Gothenburg, Sweden), and Northern Digital Inc. (Waterloo, Canada).

Other companies have developed electromagnetic motion capture systems (e.g. Polhemus, Colchester, USA) and some companies use several technologies (e.g. Ascension Technology Corporation, Burlington, USA). Electromagnetic systems use small sensors that detect the electromagnetic field generated by a source transmitter. Because electromagnetic systems do not rely on a clear line of sight between the transmitter and sensors, they are less susceptible to occlusion issues than optical tracking systems. However, metals in close proximity to electromagnetic systems cause distortions in the data, and the systems have a limited range.

Real-time motion capture is used with VEs to capture the user's position and movement, and has two main purposes. First, with stereoscopic displays in particular (discussed in more depth in the output devices section), and any type of HMD, it allows the 3D graphics to be rendered correctly, taking into account the current position of the player's head. Second, it allows the player to interact with the virtual world in an intuitive fashion, using natural movements. This could include using their body to control an avatar 'double' that is visible in the VE. This may be necessary when using a HMD, for example, in which case the player's actual body won't be visible to them. Or it could provide the player with alternative views of their actions. Alternatively, a more natural approach in many sports, particularly in situations where an athlete's movement is confined within a given space (goalkeeping, for example), is to use mo-

tion capture just to *measure* the player's movements, and use the VE only to present other players, objects etc.

In addition to capturing motion of a player's body, it is often necessary to capture motion of sports equipment such as a tennis racket, or ball, in order to evaluate a player's action (e.g., in situations where movement technique or outcome is dependent on the trajectory of the racket or flight of the ball). There are a number of sports VEs for which capturing only the motion of the equipment is sufficient (for example, V-Pong by Brunnett et al. (2006), and Tennis Space by Xu et al. (2009)). This can be advantageous because markers can be permanently attached to the equipment, minimising setup time for new users.

Another potentially important use of motion capture in sports-training VEs is as a tool for analysing the movement of the player in great detail, either in real time, or post-hoc. The standard method analysing athletes' movements is video playback (Bideau et al., 2010b). However, this has significant shortcomings, in that it relies entirely on the quality of the footage, and that it was filmed from the relevant angles. In contrast, after a player's 3D motion has been captured, it can be played back at any speed and viewed from any angle, allowing sports practitioners to study the athlete intricately in order to understand how they can improve performance. Bideau et al., for example, were able to observe previously undocumented evidence of how a goalkeeper is able to determine the destination of a thrown ball. Motion capture is not perfect in this regard; one problem is the possibility of markers being obscured, for instance, resulting in lost data, or interpolated data that is not entirely representative of the actual motion. However, these problems can be minimised by using additional capture cameras.

A further potential use of motion capture in sports VEs is to provide data for animation of virtual figures or objects. Pre-recorded motion capture is used in this way in films and games to animate computer-generated characters, because it is both faster and more realistic than manually animating a character. It is possible to edit the motion captured data to produce many variations of the same attacking or defensive moves that

would otherwise be very time consuming to create if manually animating the characters. It is also potentially possible to create models of typical behaviour into which random variability in performance is introduced. These movement variations are essential when using VE for skill acquisition since all human motion contains inherent variability.

There are other techniques that could potentially be substituted for motion capture in the above roles, and that could be easier to administer. For example, photogrammetry can be used to recover 3D structure by mathematical analysis of several photographs of an object taken from different angles. Both commercial and public-domain photogrammetry software packages are available (e.g. Eos Systems (Vancouver, Canada) PhotoModeler, Carnegie Mellon's Automatic Photo Pop-up (Hoiem et al., 2005) project). Although equipment is needed to take high-quality photographs, it is a significantly less expensive method that can produce results comparable to full motion capture suites (Chong et al., 2008). A 'depth-relief map' can also be recovered using structured-light scanning, which analyses the effects of projecting structured light into 3D surfaces (Fofi et al., 2004). A more general class of computer vision solutions have also been developed, which can analyse still or moving images to perform a task; this could be waiting for an event to occur, gathering information or controlling hardware (Moeslund and Granum, 2001).

In addition to technologies that capture human and object motion, technology for eye and gaze tracking can also be used in the context of sports training. Gaze behaviour can differ significantly between experts and novices, e.g. when tracking a ball and aiming (Vickers and Adolphe, 1997; Land and McLeod, 2000). Bad habits and unconscious actions that can affect performance can also be identified by tracking an athlete's eyes during play. There are no known examples of this in sports training VEs at present, however.

2.2.2 Hardware - Output Devices

2.2.2.1 Visual

Feedback to the user must occur in real time. Of the human senses, it is vision and touch that are the most commonly used in the current generation of VEs, and so these display modalities are explored here. The visual display technologies in use today are:

- PC monitors,
- Data Projectors, often with large display walls,
- HMDs,
- Cave Automatic Virtual Environment (CAVE).

There are many examples of each of these four visual display methods being used as part of a sports-based VE. The 'My Second Bike' project by Chung et al. (2010), was designed as a social television application to be used at home, making a television or PC monitor a vital part of the setup. Data projectors are the most common display type in this survey (figure 2.4), because they offer a larger field-of-view than a desktop monitor, and so give a more immersive experience. V-Pong by Brunnett et al. (2006), offers a good example of a large ($4 \times 3\text{m}$) rear-projected screen, used to allow the user space for movement while only being required to wear a pair of lightweight polarisation glasses to view the environment stereoscopically. HMDs are less prominent in sports VEs, primarily because they can be cumbersome to wear and can cause difficulties for users performing active tasks. They have been used in VEs in which they will not detract from the user's performance; for example a rugby VE by Brault et al. (2009) (see figure 2.5) which only requires analysis of deceptive movement, or a virtual swimming experience by Fels et al. (2005) in which the HMD is suspended with the user as part of a system of pulleys and bungees to simulate buoyancy. CAVEs are rare in sports VEs; this is most likely because they are both expensive to set up and

very space-consuming, requiring rear-projection for 4, 5 or 6 walls (including floor and ceiling). In this survey, the only project that used a CAVE was Tennis Space (Xu et al., 2009), which used a partial CAVE of 2 walls due to space restrictions.




Figure 2.5: Rugby duel simulator (Brault et al., 2009); (a) shows the virtual pitch and (b) the virtual opponent.

All of these devices—monitors, projectors, HMDs and CAVEs—can be used to present stereoscopic images, with a variety of methods for separating the left and right eye’s image when viewing the scene, including separate displays for each eye (in HMDs), red-cyan anaglyph glasses, ‘comb’ wavelength interference filters, polarization glasses, and liquid-crystal shutter glasses. There are also a number of glasses-free autostereoscopic computer monitors and televisions. Examples include Chong and Croft (2009) who used a data projector to show anaglyphic stereoscopic video in their rugby line-out VE because they believed the glasses to have better comfort, contrast, and a larger angle of head movement than polarised glasses. Xu et al. (2009) used Infitec (Infitec GmbH, Ulm, Germany) comb wavelength filters, with simultaneous projection of the two eye’s images, for their tennis VE because it eliminates the problem of ghosting even for fast movements, which is important for fast-paced sports. Simultaneous left- and right-eye projection also eliminates depth distortions in moving stimuli that can be caused by field-sequential stereo presentation. The complex problems in stereo dis-

plays are discussed in more depth later in the chapter. None of the papers reviewed directly addressed this issue.

2.2.2.2 Touch

Haptic technology provides various forms of sensing movement and providing force and tactile feedback to a user, providing them with a means of interacting with their environment in terms of touch and to enhance their immersion into the VE. The majority of haptic devices currently available on the market use a stylus grip to deliver a maximum force of between 3 N and 15 N; for example, the PHANToM product line (SensAble, Wilmington, USA), and the Omega range (Force Dimensions, Nyon, Switzerland). Such devices have been used in a sports VE (for example, in a ping pong simulator by Knörlein et al. (2007), but in general these types of device do not lend themselves to simulating sports equipment because they are generally anchored to a point in space, have a small range of movement, and cannot easily be adapted to particular requirements of different sports. Thus, a haptic device would most likely need to be custom built to allow the VE and haptic device to truly complement each other. One example is the HapStick, a virtual game of billiards with a custom haptic cue (Gourishankar et al., 2007). Also of note is the Safe Large Workspace Haptic device described by Matsuoka and Miller (1999), which was designed to provide a large workspace for sports training and rehabilitation. The end effectors of this device accommodate a variety of components and objects to grip sports equipment and support the user's body. It provides a spherical workspace of three metres diameter and delivers a maximum force of 500 N and has been used for training throwing techniques for the javelin and shot put. Support for large workspaces is an essential requirement for most ball sports training.

Vibrotactile feedback devices are a sub-category of haptic technology that involves the use of many small vibrating transducers that can be felt by the person; Alahakone and Senanayake (2009) identify rehabilitation, sport and motor learning, and infor-

mation display and gaming as relevant applications for vibrotactile technology. The use of vibrotactile feedback devices in sport is unusual, and very different to the use of general haptic devices, because they are typically not used to present naturalistic information but to signal to the user in other ways. Examples include: a motion feedback system intended to replace the auditory feedback from a coach with vibrotactile feedback cues (Van Erp et al., 2006); a system that applies vibrotactile feedback to correct movement in snowboarding (Jacobs, 2008); and a vibrotactile vest which allows the coach to direct players during football practice (Lieberman and Breazeal, 2007). Vibrotactile devices hold great potential, due to their portability and flexibility. As such they can provide additional cues in a sport VE where the need for concentrating one's attention to the environment prevents the efficient use of auditory and visual outputs. In the training process, they can be used for providing both guidance stimuli and attentional stimuli.

It is also important to consider devices that do not necessarily require any electronic components. For example, the virtual archery simulator by Göbel et al. (2010) provides the user with a real bow as part of the VE and this too, is a tactile feedback device. For sport VEs, real, physical haptic items could be of great benefit, as many sports revolve around a physical object (football, tennis ball, rugby ball, shuttlecock, etc.) that could be used to maintain links to reality.

2.2.2.3 Auditory

Audio output is another possibility. It is noted by Matsuura et al. (2006) that the sounds a system produces directly impact upon the feeling of immersion experienced by a user, and sound may provide important information for tasks such as judging how hard a ball has been hit with a bat. Using DirectSound 3D (part of the Microsoft DirectX library) to manipulate the sounds of a game of virtual hockey, Matsuura et al. achieved sound effects such as the Doppler Effect across the two networked computers being used to play the game.

2.2.2.4 Smell and Taste

The senses of smell and taste are not yet a common component of a VE and are not known to be used in sports applications. It is possible that these senses could be incorporated into sports VEs in the future to provide a more complete sensory experience to enhance immersion; the direction of a smell, for example, could provide additional cues to the sports person (Ramic-Brkic and Chalmers, 2010).

2.2.3 Software Tools

Software is needed for 3D modelling, interfacing with the hardware devices, managing the VE in real time, and providing feedback and validation. There is no specialised software for sports applications and typically VEs are built using a variety of open-source and commercial software packages. Custom software is also needed so that the full multimodal experience can be provided for the particular sports scenario. For 3D modelling, companies such as Adobe (Adobe, San Jose, USA) and Autodesk (Autodesk, San Rafael, USA) provide software specifically designed for rapid development. Such packages include Autodesk Maya and Adobe Director, and provide the user with a broad range of robust tools. A problem with commercial packages is that they are often expensive and the user is confined to the specific toolset given to them, whereas open-source software is not only free, but allows developers to modify the software for their particular project. For this reason, universities and research groups often use open-source software or develop their own. Examples include toolkits such as VR Juggler (Bierbaum et al., 2001), a suite of APIs that allow portability between different hardware solutions; eXtreme Virtual Reality (XVR) (VR Media S.r.l, Pisa, Italy; Carrozzino et al. (2005)), a development environment for virtual and augmented reality applications; CAVELib (Mechdyne Corporation, Marshalltown, USA; Mechdyne Corporation (2014)), the original API designed for use with the CAVE, which has since been commercialised; and low level interfaces such as OpenGL (Khronos

Group, Oregon, USA) and Matlab (Mathworks, Massachusetts, USA). Companies that develop specific programs for building virtual reality applications and VEs include Dassault Systemes' (Dassault Systemes, Vélizy-Villacoublay, France) 3DVIA product line (including Virtools); WorldViz (WorldViz, Santa Barbara, USA) Vizard-complete VE solutions-and extension libraries such as character creation and position tracking; EON Reality's (EON Reality, Irvine, USA) family of interactive 3D software packages; and Manageable Kinematic Motion, MKM (MKM, Golaem S.A, France).

Software for motion capture systems is usually supplied with the product. Vicon, for example, offer different software packages not just for their different hardware solutions but for different uses of motion capture. Vicon Nexus is designed for Life Sciences, Vicon Motus allows markerless capture from pre-recorded media, Vicon Blade is for film or games actors, and Vicon Polygon can generate interactive media reports for use by surgeons, researchers and physical therapists; these programs are expensive, and aimed largely at professional companies.

Haptics software usually requires an API supplied by the hardware manufacturer, such as SensAble's OpenHaptics. There are also popular open-source haptic APIs such as Chai3D (chai3d.org, 2014).

There are a large variety of software packages available for VE creation and development, but more importantly, there are few systems using one package alone. Each package or suite has different benefits, and most systems use the best available package for the specific section of a project.

2.3 The Challenges of Creating an Effective Virtual Environment for Ball Sports

As well as reviewing the common hardware and software components used in creating a ball sports VE, several challenges have been identified that need to be overcome in delivering a high fidelity training environment.

1. Do sports skills reliably transfer from VE training conditions to real-world scenarios?
2. What types of skills appear to be best suited to training in VEs?
3. Which technologies achieve the best results in specific performance measures?
4. Should stereoscopic displays be used?
5. Is high fidelity always better?
6. How, what and when should feedback be delivered to the learner?
7. When is a VE not successful in teaching motor skills?
8. What does it cost?

Each challenge is discussed in turn and the extent to which it has been overcome with reference to relevant published work.

2.3.1 Do Sports Skills Reliably Transfer from Virtual Environment Training Conditions to Real-World Scenarios?

Despite recent technological advances in VEs and support for their use in training a variety of skills (e.g., surgery, John (2008); aviation flying and landing, Hays et al. (1992); and navigation, Sebrechts (2000)), it is still not clear how skills developed through practising in a VE transfer to real-world situations. In surgical training, for example, Waxberg et al. (2004) found that whilst VE-based training can enhance skill learning above that of no training at all, it does not necessarily emulate that of actual practice. Henry (1968) proposed that the best learning experiences are those that most closely approximate the movements of the target skill and the environmental conditions of the target context (i.e., specificity of practice). For example, when participants practise under specific sensory conditions during skill acquisition, a change in these

conditions during transfer to other environments causes a significant decrement in performance both early and late in practice (Proteau and Cournoyer, 1990; Proteau and Marteniuk, 1993; Proteau et al., 1987; Tremblay and Proteau, 1998).

It has been proposed that participants develop a movement plan to optimize the sensory information present during skill acquisition and that this movement plan is specific to the information that is available during practice (Elliott et al., 1995; Khan and Franks, 2000; Khan et al., 1998; Mackrout and Proteau, 2007). Hence, if the practice conditions are changed, the movement plan previously developed is no longer appropriate for successful performance. These findings point to the importance of matching the VE as closely as possible with that of the real-world setting if transfer of learning is to be maximized. Furthermore, learning is proposed to be centred on the development of successful generalized motor programs and movement schemas (Schmidt, 1975).

These schemas are developed through exposure to four essential pieces of movement related information: (i) the movement parameters assigned to action (i.e., the force and timing of neural impulses); (ii) the initial conditions of the action (i.e., the body position, the size and weight of any object being interacted with, the environmental conditions etc.); (iii) the outcome of the action (i.e., did the action meet the intended goal); and (iv) the sensory consequences of the action (i.e., visual, auditory, proprioceptive afferent information); Schmidt and Lee (2005). Thus, if VE learning is to be successful—as measured by retention and transfer in real-world situations—the VE environment should provide opportunities to obtain all essential information required for schema development, and to a level that is as close to the real-world context as possible.

To achieve this, the VE should (i) require the learner to produce actual movements identical to those required in the real-world setting, (ii) ensure that this movement is performed with a variety of initial conditions, (iii) guarantee that the outcomes of these movements are available to learners, and (iv) attempt to provide the sensory con-

sequences of these actions. If one or more of these classes of information is omitted, or the VE does not adequately approximate the information in the real-world setting, transfer will be unlikely to occur.

Some ball-sport simulators have undergone comprehensive validation studies but this is not always the case. Where validation results are available, the evidence for skills transfer is inconclusive, and varies according to the issues discussed above. This is highlighted by two baseball case studies. Zaal and Bootsma (2011) have considered several different VE setups for the “outfielder” problem in which a baseball fielder must run to catch fly balls before the ball hits the ground. The movement of the ball, and how the outfielder might reach the right location in time, has been extensively studied both in the real world using video recordings, and in the virtual world using simulated virtual balls.

Another baseball simulator that has undergone several studies into the nature of batting tactics, motor control and behaviour of baseball players is described by Gray (2002). The players stand 3.5 m away from a large SVGA monitor holding a baseball bat that has a Polhemus FASTRAK sensor attached to the end. A virtual baseball travels horizontally towards the user from 18.5 m away (normal release location of a pitcher) but air resistance and ball rotation are ignored, leaving gravity as the only force affecting the flight path. The player swings the real bat to intercept the virtual ball. Visual feedback is given to the user to give them an indication of how well they are performing. Their results indicate that the a VE can be a useful tool for not only researching and comparing the skills of the batters themselves but also for looking into the nature of expertise.

2.3.2 What Types of Skills Appear to be Best Suited to Training in Virtual Environments?

A common goal is to improve motor control skills. Better accuracy through repeatability in performing a particular task is the most addressed skill, for example by allowing practice at aiming and throwing a ball. Anticipation and decision making skills are also well suited for training in a VE, and the level of difficulty of the task can be increased over time.

A typical example for improving ball throwing skills has been demonstrated in a rugby lineout simulator by Brook et al. (2007); Chong and Croft (2009); Chong et al. (2008) (see figure 2.6, images from Brook et al. (2007)). The rugby player sees a life-sized projected anaglyphic video of a real player jumping and the task is simply to throw the rugby ball to the recorded player. A device called ‘throw accuracy measurement’ (TAM) was created, comprising of 28 lasers forming two grids. When the lasers are intersected, the data is sent to an application running in Matlab, which determines the parabolic flight of the ball. The researchers used photogrammetry to determine whether it is a viable alternative to conventional motion capture systems. Not only did the photogrammetry system meet the accuracy requirements of the researchers, it was considerably cheaper and more portable. Tests conducted to measure the accuracy of players from different levels of play (experienced and novice) show that on average throws to closer targets were more accurate than throws to farther targets, and that inexperienced players made more than twice as many errors as experienced players.

Goalkeeping is another example where in a training session the keeper is often required to fend off a continuous stream of balls fired at him/her. Bolte et al. (2010) created an effective system for a handball goalkeeper, shown in figure 2.7. The main elements of the system are a projector or HMD (Carl Zeiss Cinemizer Plus–stereo, by Carl Zeiss AG, Oberkochen, Germany) to show a clip of an attacker attempting to score, a full-sized goal with a uniformly coloured rectangle in each corner, and a




Figure 2.6: Rugby lineout simulator (Brook et al., 2007); (a) and (b) show the user throwing the tracked ball through the TAM device.

webcam to capture the keeper. The keeper is shown a clip of another player attacking, and moves to catch the ball as normal. The webcam uses computer vision techniques to watch the coloured rectangles in each corner of the goal; when the keeper moves to block the attempt and covers one or more of the patches with his hands or feet, the webcam detects this. At the end, the results are examined and the player is given a score. This extremely cheap and simple system is now used by the German Handball Federation to train their youth national teams and to help find new young talent, though the effectiveness of the system has not been reported since.

2.3.3 Which Technologies Achieve the Best Results in Specific Performance Measures?

The technology deployed does influence performance measurements such as the success rates of completing a task e.g. catching a ball. Currently the choice of visualization technology is the dominant factor, and the necessity to wear glasses, HMDs and/or tracking equipment invariably provides an unwanted constraint on the user. Utilisation of other technologies such as haptics is also starting to appear and will likely have a growing influence on performance measures.

A system to investigate the techniques of catching fly balls (Fink et al., 2009) uses a



Figure 2.7: Virtual handball; note the blue detection patches in the four corners of the goal; picture courtesy of Bolte et al. (2010).

HMD (Cybermind Visette-Pro; Cybermind, Maastricht, The Netherlands) and a baseball glove worn on the left hand, both tracked using the IS-900 hybrid inertial/ultrasonic tracker by InterSense (Intersense, Billerica, USA). There is a 12 m² area in which the users can freely move. The user can see a basic virtual scene of the ground and sky, and are tasked with catching virtual baseballs, which are fired into the air from 36 m away. The balls then follow either a parabolic path according to normal gravity conditions and no air resistance, or a perturbed path where after reaching the highest point, the ball drops vertically at a constant velocity into one of the four quadrants surrounding the player.

The authors decided that as they were not studying the actual catching of the ball but the act of moving to the correct location, participants needed only to move their head or glove within 1 m of the ball to “catch” it. Results of a study of 8 experienced

male and 4 female college baseball and softball players showed that participants caught between 44–55% of the balls (depending on the conditions of the path), which is much lower than might be expected in the same situation in the real world. The authors suggest that the lower percentage of successful catches may be due to the weight and limited field-of-view of the HMD. The CAVE-based study for fly-ball catching (Zaal and Michaels, 2003) also presented problems as the available space was too small for any tasks requiring movement; the system also lacked ceiling projection (and so lessened the sense of immersion), and suffered from a time delay (latency).

A CAVE was also used for Tennis Space (Xu et al., 2009) where attempts were made to overcome space restrictions without requiring the user to wear a HMD. It was decided that one sidewall was necessary but that using front, left, and right walls would be too restrictive, and that a floor would be more immersive than a ceiling panel. It was not possible to rear-project onto the walls, or to raise the floor to project from below. The solution to this was that the front and right wall were projected from above the player, casting a shadow behind them rather than in front, and the floor was simply covered with carpet, carefully matched to the colour of the grass projected onto the screens. A virtual tennis court and virtual opponent (VO) were created using OpenGL Performer and CAVELib; the VO was animated using motion-captured data of a real tennis player. A hybrid tracking system was created, which used both an ultrasonic tracking device to track sensors on the stereoscopic glasses (Infitec) and on the tennis racket and four cameras to track the body of the player. The authors concluded that the projectors were not bright enough at 12000 ANSI Lumens, and the size of the room in which the system was constructed interfered with the design: in a larger room a floor projector would have been used instead of a carpet and the two wall panels would have been made taller to cover as much of the peripheral vision of the player as possible.

A VE does offer a flexible environment for gathering and analysing performance data. This is demonstrated, for example, by the handball skills training environment by Bideau et al. (2003). The system consists of the motion-captured participant view-

ing the stereoscopic VE through shutter glasses; the handball goalkeepers see a VO running at them with a virtual ball, and they have to prevent the VO from throwing the ball into the net and scoring. The project initially involved observing how goalkeepers reacted to the VO compared with how they reacted to human opponents. It was found that the use of motion captured attack moves applied to the VOs provoked real defensive movement from the goalkeeper (Bideau et al., 2004a). The visual information the handball goalkeepers use (Bideau et al., 2004b) was also analysed, and it was found that the use of a VE allowed the modifying of specific individual parameters, producing an excellent testing environment.

Follow-on work examined the participating goalkeeper's behaviour in judgement and motor tasks (Vignais et al., 2009b), differences in their performance for different qualities of graphical detail (Vignais et al., 2009a, 2010b) and visual conditions (Vignais et al., 2010a). As well as handball, the research group investigated deceptive movement in "sporting duels" between rugby players (Brault et al., 2009, 2010); shown in figure 2.8). Here the user wears an HMD (Cybermind Visette Pro with dual LCDs); the users' avatar was controlled using the left and right arrow buttons on a Microsoft Sidewinder gamepad. The VO appears in front of the user, running towards them with a rugby ball. The user has to decide which direction they believed the VO is going to go, while the VO used deceptive motions to trick the user by deliberately shifting their centre of mass to confuse and disorient them.

Apart from vibrotactile effects found in devices such as the Wii Remote, haptic feedback has also been used successfully in sports-themed VEs. One of the first examples to successfully exploit haptics is in a Ping-Pong game described by Knörlein et al. (2007) where virtual bats were co-located with two SensAble PHANToM 1.5 three-degrees-of-freedom force-feedback devices so that players could feel the impact of the simulated ball on the bat. Although this improved hand-eye coordination, the limited workspace of the PHANToM and lack of six-degrees-of-freedom force feedback were limiting constraints. The potential use of haptic guidance (i.e. non-natural feedback

about performance, as discussed earlier) in training a dynamic motor skill (including in sports) has been studied in more detail by Huegel and O'Malley (2010). Their findings show that haptic guidance is useful to trainees when beginning to understand the task but should be “progressively removed to avoid possible negative dependence”, and al-




Figure 2.8: “Sporting duels” between rugby players; real data is captured to animate virtual opponents. The user wears an HMD and indicates the direction they believe the opponent will move on a gamepad (Brault et al., 2009).

low the trainee to become reliant on their own skills. When all force is disabled, the haptics device can still be used for tracking and evaluation of the player's performance.

2.3.4 Should Stereoscopic Displays be Used?

On first inspection, it may seem obvious that VEs for ball sports should present stereoscopic depth information because binocular disparity is available in natural viewing, and results in a vivid sense of three-dimensional structure. The majority of the projects identified in this survey do indeed use stereoscopy. There are, however, practical, perceptual and ergonomic drawbacks of stereoscopic presentation, and these can be particularly acute in environments like sports VEs, which frequently need to present rapid motion, and large ranges of depicted depth.

First, stereoscopic VEs are generally significantly more expensive and complex than otherwise equivalent non-stereo systems, because of the specialised technology required. Clearly these extra 'costs' must confer sufficient benefit in terms of the system's effectiveness to be worthwhile.

Second, although in entertainment applications the terms 'stereoscopic' and '3D' are used interchangeably, a person's perception of depth in fact relies on a number of depth 'cues', the majority of which are monocular (Howard, 2002). Depending on the task, perfectly adequate 3D percepts may be evoked by non-stereo displays. Psychophysical research suggests that for perception and for visuo-motor control, the brain combines signals from all available cues (visual cues such as binocular cues, texture/perspective, motion) weighted according to their relative reliabilities (Howard and Rogers, 2002; Hillis et al., 2004; Jacobs, 1999; Keefe et al., 2011; Knill, 2005; Knill and Saunders, 2003). Because the reliability of depth from binocular disparity falls off rapidly with viewing distance, even at near distances, depth information from monocular texture/perspective cues can sometimes be as reliable (or more) than binocular cues (Hillis et al., 2004; Jacobs, 1999). Motion Parallax (which requires head tracking) can itself provide a powerful percept of depth, and is likely to be important in many sport-

ing situations, where the player typically moves. The importance of presenting stereo depth cues is therefore likely to depend on the specific details of the environment being simulated.

Third, viewing conventional stereoscopic displays often results in fatigue and discomfort, difficulty fusing stereo image pairs, and poor perception of stereo depth (Akeley et al., 2004; Emoto et al., 2005; Hoffman et al., 2008; Ukai, 2007; Ukai and Howarth, 2008; Wann and Mon-Williams, 1997; Watt et al., 2005b,a; Yano et al., 2004) all of which are particularly unwanted in a sports trainer or simulator. These undesirable effects have all been shown to result, at least in part, from the unnatural stimulus to the eyes' focusing response in stereo displays. Because the two eyes' images are presented on a fixed screen, the observer must converge his or her eyes at the portrayed depth, while focusing (accommodating) to the screen surface, resulting in a vergence-accommodation conflict. The required decoupling of accommodation and vergence responses is effortful, and sometimes impossible, because the systems are synergistically linked (Fincham and Walton, 1957; Martens and Ogle, 1959). Successful solutions to this problem have been developed (Akeley et al., 2004; Liu et al., 2008; Love et al., 2009; MacKenzie et al., 2010) but most require the head to be fixed relative to the display (as in an HMD) and even then are not yet practical for a wearable display (with the exception of Love et al. (2009)). Moreover, none are yet commercially available. At present, therefore, it is important to ensure that stereoscopic content remains inside a "comfort zone" around the screen distance (Shibata et al., 2011), and that the distance to the screen is as closely matched as possible to the depth in the portrayed scene. This may be difficult to achieve for stimuli such as balls flying out of the screen towards the viewer.

Fourth, stereoscopic images appear very distorted when viewed from the wrong place, and appear to move if the graphics rendering is not updated to reflect the current eye position. Although there can be only one geometrically correct viewing position for any image, perceived 3D structure from conventional '2D' pictures is surprisingly

consistent when viewing from other positions (Vishwanath et al., 2005). In stereo images, however, depth can appear very distorted, and if the observer moves without compensation in the images on the screen, the 3D structure of the scene appears to distort dramatically, something that would be highly undesirable in a sports VE (Banks et al., 2012). You can easily see this effect by moving your head side-to-side while viewing a stereogram. As noted earlier, this problem can be solved by tracking the head position and updating the two eyes images correctly, in real time. Note, however, that the orientation of the head must also be tracked and corrected for, or the player may tilt his or her head to one side, requiring vertical vergence eye movements (moving one eye up, and the other down, relative to the head) to fuse the stereoscopic images. Such eye movements are rare in natural viewing and are known to cause significant discomfort (Kane et al., 2012).

Fifth, many current stereoscopic display systems present the left and right eye's images field-sequentially, which can result in perceivable judder in moving stimuli, especially for rapid motion. Attention is frequently paid to reducing perceived flicker, by using presentation rates (at each eye) that exceed the visual systems flicker fusion rate (~ 60 Hz, Hoffman et al. (2011)). However, motion judder is in part caused by the insertion of blank frames in each eye's 'image stream' - the left eye's image is blank while the right eye's image is displayed, and vice versa - and remains perceivable at presentation rates that are high enough to avoid flicker (MacKenzie et al., 2010). Field-sequential stereoscopic presentation can also result in distortions in perceived depth, in which the perceived trajectory of laterally moving objects differs significantly from that in the original scene. The two eyes' images are captured simultaneously (by pairs of cameras), but presented at different points in time. Thus, for a given left-eye frame, for instance, there is uncertainty about whether the brain should compute binocular disparities by comparing it to the previous, or next, right-eye frame. Because the visual system integrates over time, this leads to average biases in the estimate of depth from binocular disparity (Hoffman et al., 2011). Like motion judder, this problem

is dependent on the speed of the moving stimulus, and so is likely to be particularly problematic in sports simulators. Note that in practice, these problems are not limited to shutter glasses systems, in which field-sequential presentation is an inherent feature, because systems based on comb wavelength filters or polarisation also often use field-sequential presentation so that only one projector is required.

Clearly stereoscopic presentation can, in some circumstances, significantly improve depth perception, and the viewer's sense of immersion in a scene. However, given the numerous problems with current display technologies, and the fact that many of these are particularly acute in situations with large changes in stereoscopic depth, and moving objects, the use of stereoscopic presentation in virtual sports environments needs to be considered carefully, in light of the particular features and requirements of the depicted scene (see Banks et al. (2012) for a review).

2.3.5 Is High Fidelity Always Better?

The level of realism achieved in different VE systems can vary enormously and there are potentially important trade-offs to be made between what can be achieved in a practical, affordable system with current technology, and what is desirable for maximum effect. An important general distinction can be made between the extent to which the sports VE should depict the appearance of the real world (perceptual fidelity) and the extent to which it should behave like the real world when a participant interacts with it (functional fidelity).

One highly relevant area is in deciding the level of graphical realism required. Graphical realism is constantly improving, mainly due to the popularity of CGI films and video games, but photorealistic rendering is computationally expensive. It might initially seem that the more like the real world the VE appears, then the better will be the athlete's performance during skills acquisition exercises. However, even simple stimuli such as animated point-light figures can be effective in presenting movements

of an opponent (Vignais et al., 2009a, 2010b), suggesting that the level of rendering detail is not necessarily a significant factor in the success of acquiring a new skill.

A key aspect of the visuo-motor skills involved in ball sports is that they require perception of the movement of players and objects in three dimensions. Thus, from a purely perceptual standpoint, it may be possible to present information about 3D structure and motion in the scene in a highly simplified manner, and still achieve training effectiveness. The challenge here is that identifying the information used in real-world tasks is far from straightforward, and there is typically no simple answer. Consider catching a ball, for example. This requires judgements about the trajectory, speed, and ‘time-to-contact’ of an approaching object. A large body of work in psychophysics has shown that this can be achieved using a combination of information from visual direction, binocular disparity, inter-ocular velocity differences, changing size, and to some extent even knowledge of an object’s size (Rushton and Wann, 1999; Regan and Gray, 2000; López-Moliner et al., 2007).

Recent studies are consistent with the idea, described earlier, that the brain integrates information from all available sources of information, weighted according to their reliability (Lages, 2006). Critically, this reliability changes dynamically, depending on geometrical factors (Hillis et al., 2004; Lages, 2006). Thus, the likelihood is that there is no fixed answer to the question of how much ‘weight’ a given cue receives generally, or even in a particular task. This means that while individual cues, in simplified environments, may evoke a compelling percept of 3D motion, the brain is not necessarily relying on the same information it would do in a detailed real scene. Details of any optimised rendering techniques used (removing certain depth cues) is typically omitted from the papers reviewed. However, it appears from the sample images provided that the graphics are rendered with correct perspective and use standard shading and lighting techniques (the use of shadows as an extra cue for depth perception is not often used, and no example has been found where shadows are applied to the player avatars). Ideally, the implications of such approaches for training need to be

established empirically, so that rendering can be simplified without loss of effectiveness.

The above discussion considers how graphical realism may not be critical for producing veridical perception of a three dimensional scene. It may be important, however, in creating the sense of immersion or “Presence” in a VE. Presence is a much-debated concept, with many different definitions having been proposed (Slater et al., 1994; Schuemie et al., 2001). The current definition of presence in a VE is “the feeling of being there”. The standard method of measuring presence is a questionnaire, although biomechanical analysis has been investigated (Bideau et al., 2003). Both a pre-experimental questionnaire and a post-experimental questionnaire are used. Video footage of real opponents can be effective in a sports scenario, for example as used by Chong and Croft (2009), and virtual characters or opponents assist in improving the realism of the environment or to get the user to interact, such as in an attacking move Rusdorf et al. (2007). However, virtual characters and environments do have a wider range of functions available, mainly because once a video is made it is unchangeable, whereas a VE can react to the user and change the game as it is being played. Video footage is ideal if the virtual characters or environment are required to execute the same simple procedure every time, but this will not be an acceptable method of engaging the player. Presence could be an excellent benchmark for VEs, but will be unusable until a definition is agreed upon and a standard method of measurement is established. Moreover, the importance of presence in training effectiveness needs to be evaluated empirically for different situations and tasks.

The use of virtual characters leads to potential issues with the “uncanny valley”, described by Mori (1970). This refers to the discomfort, and sometimes fear, felt by some people regarding a robotic or virtual human with aesthetic features very similar to a real human (Seyama and Nagayama, 2007). In a VE, if any aspect of the scene is overly distracting or unpleasant, it is likely to have a negative effect on the user.

As well as the properties of the user interface, the degree of functional fidelity

of the modelled world must be considered. A high functional fidelity simulator will typically use physics models to provide the most realistic simulation of an event that is possible, which is a challenge for real time response. For sports-based VEs, this applies specifically to the motion of virtual objects such as a ball or projectile (shuttlecock, arrow, etc.) when travelling through the air or rolling along the ground. The real physics parameters (such as surface properties of a ball) are needed for the mathematical model, but can be difficult to obtain without the involvement of an industry partner. Often when companies do undertake such research, it is kept private to avoid aiding competitors, making it difficult for academic groups to acquire the appropriate information, which increases the difficulty of building an accurate model simulation.

Academic groups are often forced to either search for existing free-to-use parameters from other research groups or to perform in-house experiments to derive or acquire the required parameters. Further, due to high computational demands, real-time simulations based on accurate physics models are hard to achieve. This leads to an important trade off: the high complexity of a mathematical model produces more accurate results, but involves more parameters and requires greater computational power—often too much for a real-time solution. A compromise must be reached whereby limitations and assumptions simplify the modelling of the problem, but also reduce the accuracy; the chosen solution depends on the resources available to the group and the complexity of the model. As with the example of depth information, perceptual and motor experiments can in principle be carried out to determine which simplifications to the physics models can be made without incurring costs in training effectiveness, but this is likely to be a difficult and time-consuming process.

2.3.6 How, What and When Should Feedback be Delivered to the Learner?

In order to learn a skill, feedback can be delivered to the user in a number of ways. Eaves et al. (2011) investigated the use of VEs for teaching real-world motor skills for dance and sports. They found that the ability to direct the learner's attention to the key anatomical features of a to-be-learned action was important. Ruffaldi et al. (2011) discuss different types of feedback appropriate for sports training in a VE. Rowing is used to illustrate their work, but the points raised are equally relevant to ball sports. The two main categories of feedback are:

- Informative feedback, which involves supplying the user with information and statistics regarding their performance (e.g. knowledge of results such as a score);
- Guidance feedback, which directs the user about how to perform the next action (e.g. knowledge of performance such as body movements).

Salmoni et al. (1984) provide a classic review of the effects of information feedback on motor learning. In both cases the feedback can be delivered in various ways: during (concurrent feedback), immediately following (terminal feedback) or some period after the completed skill (delayed feedback). In reviewing these various feedback deliveries, Salmoni et al. proposed the guidance hypothesis. This states that feedback essentially guides the performer to the correct movement pattern and increases skill acquisition. However, if the learner receives feedback too frequently they may develop a dependency on it, and thus ignore the processing of important intrinsic sensory feedback required for error detection and correction. Consequently, when augmented feedback is not available, performance suffers, as the learner has come to rely upon it to produce the required skill effectively. Thus avoiding negative training can be helped by progressively reducing the feedback presentations across the learning cycle and thus removing the dependency effect. Ruffaldi et al. (2011) note that the individual methods

of feedback are rarely considered carefully in designing a VE, although they ought to be.

Methods of reducing the amount of required learning time can be employed by using specific forms of feedback and analysis, known as “accelerators”. Accelerators can be categorised into augmentation (Ruffaldi states: “the ways the training experience is enriched with respect to the real situation”)—information, correction and task enrichment; simplification (“it tries to make the user perform the task in an easier way with respect to the way the task is performed in the real world”) and variability (“changes in the training environment aimed at forcing the user to leave a stabilized but novice behaviour and move toward a transitory unstable expert behaviour”).

An example of informative feedback can be found in the golf swing trainer described by Kelly et al. (2010). Forty male golfers were recorded using a Vicon motion capture system while attempting to hit the ball into a net 3 m away. The players were split into groups of different skill levels; the “expert” level players’ movements were averaged into a single movement, which was applied to the “coach” avatar. Users’ motions were then captured and displayed as an “apprentice” avatar next to the coach, allowing them to compare their movements with those of the coach in great detail. The user had full control over the zoom, angle, playback speed and positioning of the coach; they also had access to graphs comparing the different “angle[s] of interest” with the coach. It is also interesting to note the findings of Chiviakowsky and Wulf (2007), however, who investigated the motivational aspects of feedback and demonstrated that learning is facilitated if feedback is provided after good rather than poor trials. Visual and haptic cues can also be used to provide guidance. For example, the optimum trajectory that a ball should follow to reach a desired target could be displayed in the VE using appropriate graphical glyphs. None of the projects examined in this survey currently do this, however.

2.3.7 When is a Virtual Environment Not Successful in Teaching Motor Skills?

A danger with any VE that aims to train a particular skill is that it trains the skill incorrectly and results in a negative efficacy. This is just as much a danger in a sports trainer, although the side effects may not be as catastrophic as negative efficacy in a flight simulator! Li et al. (2010) discuss negative efficacy in the context of using haptics to improve performance of manual control tasks. Their experiment, based around a target-hitting control task, indicated that the acquisition of motor skill is a complex phenomenon that is not aided with haptic guidance during training for that particular task. Although as noted above, progressive use of haptics does offer benefits. Results such as these must be considered when designing and building a sports VE.

Another problem is time delays, or latency, particularly for networked VEs or VEs for fast-paced sports such as table tennis. Here, even a small delay can make the game highly unnatural, and even unplayable. Latency also reduces the feeling of immersion in a VE, one consequence of which is also to contribute towards negative efficacy of the training experience. Consider the V-Pong system (Brunnett et al., 2006; Rusdorf et al., 2007), shown in figure 2.9. Two tiled projection screens form a stereoscopic image behind two 2 m × 3 m wall panels. The player sees the table directly in front of him as though he is playing a real game, with a VO standing at the other end of the table. Restrictions are placed on the VO to ensure it is beatable; a maximum speed for its movements is set and random noise is added to its movements to stop a perfect performance. The biggest issue with this system was that the speed of a game of table tennis is much higher than the system can cope with. Latency is a problem with many virtual reality systems, but with table tennis it was far more obvious. To tackle this problem, the team devised a prediction system which predicts the next location of the tracked racket and glasses from the speed and direction they are currently moving in,

accounting for the approximate latency of 50 ms. The prediction system is adjustable, allowing for changes to be made for different people who move at different speeds.



Figure 2.9: V-Pong table tennis; picture courtesy of Brunnett et al. (2006).

The need for high speed tracking was also demonstrated in a two-player table tennis game (Li et al., 2010), seen in figure 2.10, in which the players are immersed into three screens in the shape of 3/8ths of an octagon. Each player wears circularly polarised glasses, which show the screen in stereo and are motion tracked, and they each hold a tracked wand to represent their rackets. To model the ball, simplified physics laws were applied for collision detection and movement. The system was developed largely to test a new high-speed wireless tracking system, which was found to be responsive to the user. The system was successful in displaying the opposing player's movements through the on-screen avatar, although the virtual scene itself was very simplistic using basic colours and avatars with rigid bodies. It is noted that for future work, the authors would like to not only create more realistic avatars, but also improve the physics and utilise full-body tracking.

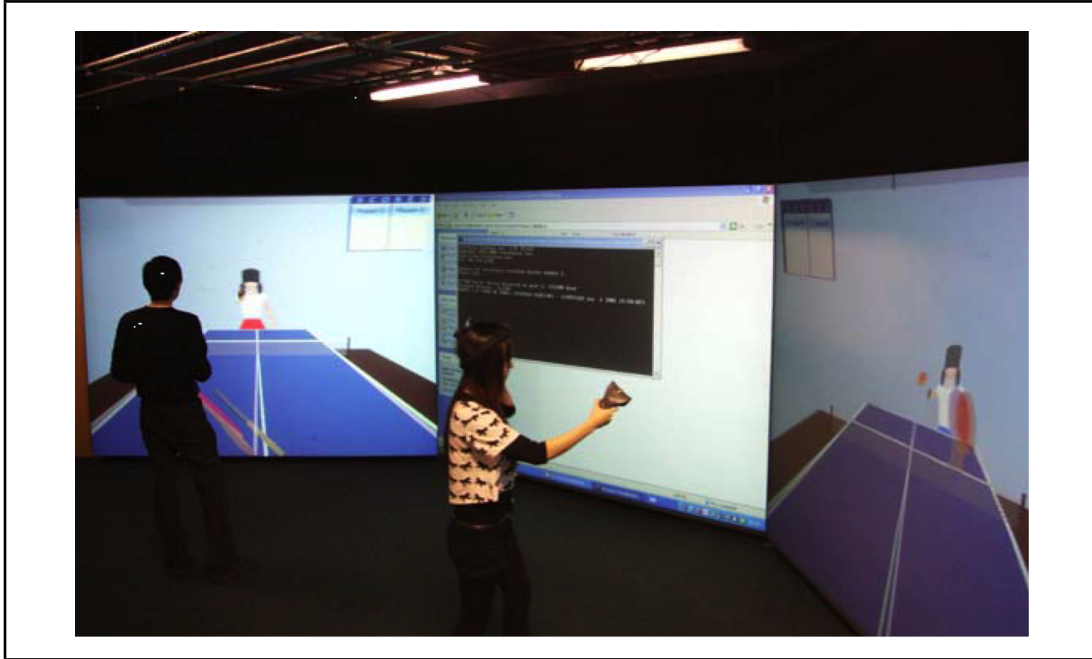


Figure 2.10: Virtual table tennis for two players; picture courtesy of Li et al. (2010).

Of course, for any VE to successfully teach motor skills, or anything else, it must minimise side effects such as latency because they can also cause cybersickness (Wloka, 1995; Potel, 1998), due to the conflict between vestibular signals to the brain and the (delayed) visual motion information. Cybersickness does not affect all VE users, and it is believed that the brain may be able to adapt and become tolerant to sensory conflict after prolonged exposure to situations in which it is common (Potel, 1998). Nonetheless, validation studies are required to confirm that training effectiveness is not affected.

2.3.8 What Does it Cost?

A further problem involved with creating an effective sports-themed VE is the cost. Costs associated with all kinds of VEs include the cost of software, hardware and maintenance, and validation. Unless the software is developed entirely using open-source packages, it is likely to be expensive. Whilst some VEs discussed in this chapter use only readily available hardware, some use specialist hardware such as HMDs, CAVE

setups and motion capture systems, which are not only very expensive to purchase, but can also be expensive to maintain. The projects that are described above range in cost from around €10,000 to well over €250,000 (if a CAVE is used, for example). It should be noted, however, that there are cheaper VE installations with good functionality available.

Robustness of equipment should also be considered alongside cost, as some equipment may be subject to more stress than normal in active sports VEs. An appropriate investment must be made, however, or otherwise the training benefits achieved will be marginal at best.

2.4 Summary

Training of sports-related skills is a fast emerging application of VEs and there are many different options for software and hardware, each with their own benefits and problems. There is no single solution and the choice of the equipment deployed is largely based on what suits each specific sport best. In this chapter, the different concepts and technologies involved in current VEs in the context of ball sports have been reviewed and the key research and development issues that need to be considered by researchers and potential users of this technology have been identified. In particular, eight different challenges/issues have been categorised and the extent to which they have been overcome has been discussed with reference to the current state-of-the-art.

The cost range of the systems in current use is large but it is not always the most expensive systems that produce the best results. For example, although most systems have invested in display equipment that will offer stereoscopic viewing, the practical, perceptual and ergonomic problems of this technology means that its use should be carefully considered. Indeed, the amount of both perceptual and functional fidelity required to effectively train a task is still unclear and in some cases a low fidelity, inexpensive system can perform well.

Another common problem is how to represent a potentially large playing area (e.g. a football field) within the physical constraints of the VE? Any constraints that take you further away from the real training environment are going to impact on the functional fidelity of the VE. Removing the inherent constraints of the hardware deployed is an open research topic. Even when a good representation of the sport can be constructed within the VE, there remains a danger of delivering negative efficacy in the training task. Latency is one contributory factor in negative efficacy, but there are other potential pitfalls such as how to apply haptic cues appropriately. This is another area where more research is needed. On the question of feedback, then there are successful examples of providing informative feedback in a ball sports VE. The development of guidance feedback is far less advanced, however.

One advantage of using a VE should be that it provides the flexibility to conveniently allow learners to practise a wide range and number of skills. This is fundamental to improving performance. However this requirement has not yet been fully addressed, as most of the reviewed examples tend to implement just a small number of skill scenarios. This will be particularly important to address for commercial products.

As the domain of sport-themed VEs matures, then more validation studies will be possible. Results of the currently published validations are inconclusive, with evidence being presented that in some cases there is skill transfer but in other cases there is not. There are enough results, however, to suggest that the improvement of motor control skills in ball sports is possible. Anticipation and decision making skills are also well suited for training in a VE, and the level of difficulty of the task can be increased over time. It is predicted that this area will grow quickly in the next few years with more and more examples coming out of research laboratories and being used by professional sports associations.

The issues and solutions identified and discussed in this chapter will be invaluable to the goal of producing an effective training environment for rugby passing skills.

Chapter 3

VE Testbed

To investigate the hypothesis “A virtual environment can be used to effectively train ball passing skills in rugby’, a virtual environment must be built for the purpose of training passing skills.

The previous chapter highlighted several useful points for building a ball-sports VE. In a survey of similar systems, the most popular motion capture technology was found to be optical motion capture. This is because it provides high-performance capture with devices able to capture at high frame rates, and it is easy to mark the user or an object used in the VE with retroreflective markers. The most popular display technology was the data projector, as it is a simple way to display a virtual scene in front of the user, and can offer stereoscopic display when the projector is able to display with a 120 Hz refresh rate. A VE should be designed carefully to ensure it is easy and natural for the user to interact with, and so haptics and tangible interfaces are also popular. A further design consideration for VEs, particularly in ball-sports VEs, is latency. Objects moving at high speed will make any latency more apparent by presenting motion judder. Any latency in the VE may also cause the viewer to suffer from cybersickness.

A testbed system was designed and developed to prepare the basic functionality of the system. This chapter describes the process of creating the testbed system. It

has been divided into sections based on each portion of the research rather than in a specific timeline, as most of this work was occurring concurrently.

The first section describes the planning stage, and shows how the research undertaken was initially shaped by the designs created; the analysis in chapter 2 was used to select the hardware and software for the VE testbed. This leads into the task of building the system, starting with the core functionality for the ball tracking and movement; then the construction of the scene (originally using the XVR authoring package) will be detailed piece-by-piece, followed by the additional program elements—the user interface and control options, and how the scene was made interactive. Finally the evolution of the physics engine from basic straight-line equations to a complex system involving the effects of the weather is described.

3.1 System Design

The first decision to be taken about the system was what the focus of the VE training should be, not only in terms of the skill to be trained but also the subject who would be using the system and what reason there is for it to be used. There were a number of basic skills in a game of rugby union which had potential to be evolved into a virtual training drill: passing, lineouts, kicking and even the scrum were considered. Users of the VE could vary from professional rugby players of local or international teams, Under-19's (also local or international) and junior players. Reasons for use extend beyond training in a basic repetitive drill to additional functionality such as using statistics about the player to adjust posture and technique and even assessing the potential of younger players (talent scouting).

Because this research is being undertaken on behalf of Rygbi Innovations Cyf., these options were discussed at length to ensure it met the requirements of the company. Discussions with a team coach from the Colwyn Bay rugby club lead to the passing and lineout skills becoming the focus of the research: passing skills are often

overlooked as an essential skill to the game and lineouts require an entire team to practise in the traditional method. Virtual lineout training had already been the focus of a research project (Brook et al., 2007; Chong et al., 2008; Chong and Croft, 2009), so the task of creating a VE for training passing skills was adopted as this would be both novel and complex to achieve a high fidelity.

A plan for the system was to create a virtual pitch environment with a target to aim the ball at in a pass move. A major benefit of using a VE to provide a training platform for training any skill is that users can be presented with an environment that might not be easily accessed on a regular basis. This could be providing different weather conditions (wind and rain), or environmental conditions such as heat and air pressure effects on the ball's movement. It would be difficult to present these conditions to users, short of listing the effects before they enter the environment. Presenting wind direction and strength may be possible, to an extent, by introducing flags or even a windsock into the background of the environment. This would provide users with a cue that they should be able to associate with the real world.

Some requirements and preferences for the software used to create the VE were identified:

- 3D content authoring to create the virtual scene;
- stereoscopic rendering to explore whether the use of stereo improves the VE;
- free to use; to allow for uses with a limited budget;
- ability to develop on a Windows 7 64-bit PC (the PC provided for the development work);
- software that can accept a variety of input devices to allow the upgrading of the tracking hardware;
- provide tools to produce a program which can be run easily by non-specialists;

- the program produced should not be resource intensive and, if possible, run platform-independently, since the specifications of the future running system are unknown.

3.1.1 Hardware

A small number of sports VEs (Brook et al., 2007; Chong et al., 2008; Chong and Croft, 2009; Ahn et al., 2006; Mestre et al., 2011; Göbel et al., 2010; Gourishankar et al., 2007; Chung et al., 2010) were found in which the user holds the actual apparatus used in the sport to improve the fidelity of the virtual training environment. In the case of a passing task, it would be very difficult to undertake the training without holding a real ball, as many of the factors contributing to the player's sense of the pass revolve around the ball itself. The weight and size of the ball coupled with the player's hold, grip and the amount of force put into the pass are critical. Without having a real ball present, this information would not be available and would make an accurate pass very difficult to perform and a drill more likely to be detrimental than useful.

The VE would be designed to use a variety of different display and tracking technologies so that it can be customised according to the budget available and the level of fidelity required. There is a trade-off between these two requirements, as discussed in chapter 2. The hardware requirements would be a real ball, a method of tracking the ball's motion, a graphics engine (PC) to process the motion data and a device to display the VE to the user. The display, as discussed in chapter 2, could be screen-based or a HMD; the screen-based approach was chosen because it would be less restrictive to the user. For development, the Personal Space Tracker (PST) (Personal Space Technologies B.V., Amsterdam, NL; see figure 3.1a) was used, with the intention of using a device that provides higher-quality tracking for the final product. The PST is a compact optical tracking solution that utilises infra-red light to highlight the movement of retroreflective markers placed on an object (see figure 3.1b). Note that the 50 Hz (PST-55) capture-rate was the only model available when purchased; newer models




Figure 3.1: (a) The Personal Space Tracker optical tracking device. (b) The retroreflective markers as seen by the infra-red cameras.

have since been released. The PST was selected because it offers high-quality tracking at a relatively low price; all model processing and kinematic solving is done within the unit itself, and it is more portable than many of the tracking devices on the market thanks to the small size of the unit.

Objects to be tracked have retroreflective markers attached to their surface (see figure 3.2a), then PS Track, the software that processes the PST data, can be trained to recognise the individual object. During the training, the user slowly rotates the object

in front of the camera until all points have been found; PS Track then determines the dimensional relationship between each marker and its neighbours to create a unique map of the object (see figure 3.2b). The point in the centre of the marker map is returned as a single point in the calculated centre of the object.

Once a model has been recorded and analysed by PS Track, several objects (up to 20 individual objects per tracker) can be simultaneously managed in one scene without issue (see figure 3.2c and figure 3.2d).

To meet the requirement of building a system for which the input devices can be altered with minimal editing of the program, an interfacing program would be required. A number of interfacing systems were evaluated, including:

- Conduit by Mechdyne;
- Trackd by Mechdyne;
- VR Juggler by Iowa State University's Virtual Reality Applications Center;
- VRPN by the Department of Computer Science at the University of North Carolina at Chapel Hill; and
- MiddleVR by I'm in VR.

Conduit and Trackd were discounted over the cost of a software license, leaving VR Juggler and Virtual Reality Peripheral Network (VRPN). MiddleVR is able to configure complex systems, and would not be used for the initial testbed system as that level of complexity is not yet required. Despite VR Juggler providing a varied suite of programs to handle a variety of aspects in the management of a VE, VRPN was chosen for the in-built compatibility with PS Track and for the large range of compatible devices it boasts. A table containing more detailed evaluation can be seen in appendix 7.1.




Figure 3.2: (a) A rugby ball covered in retroreflective markers. (b) The model of the rugby ball in PS Track; each point represents a retroreflective marker affixed to the surface of the ball. (c) Three tracked objects in the PS Track software. (d) The three tracked objects in real space.

Software packages were then evaluated in a similar format (see appendix 7.1). The packages were chosen primarily for their content-authoring ability. Additional requirements and features were considered once the list had been created:

- 3DVIA Virtools by Dassault Systemes;
- Avango by Fraunhofer IAIS;
- CAVELib by Mechdyne;
- EON Suite by EON Reality;
- Maverik by AIG at Manchester University;
- Ogre3D, a freeware program;
- OpenSceneGraph, an open source project;
- Panda 3D by Carnegie Mellon University;
- Syzygy by the University of Illinois at Urbana-Champaign;
- Unity 3D by Unity Technologies;
- Vizard Virtual Reality Toolkit by WorldViz; and
- XVR by VR Media S.r.l.

Several options are immediately discarded because of incompatibility with Windows 7, the cost of the software being beyond the scope of the research and lack of stereo content support. The remaining contenders were Avango, OpenSceneGraph, Panda 3D, Unity 3D and XVR. After consideration, XVR was chosen for its in-built support of VRPN. However, the ability to use other software packages where appropriate remained a possibility and the software components were developed in a modular format to facilitate this. For example, the experiments described in chapter 4 made use of a Virtools port of the VE.

A significant feature of the VE software was the physics engine to calculate the movement of the virtual ball. Despite many of the reviewed software packages containing inbuilt (or support for external) physics engines, this was not considered a requirement as it would be manually developed as part of the research.

To summarise the plan for the research: the prototype testbed environment consists of a real ball to be tracked and aimed at a virtual target on a screen before the user; a combination of XVR and VRPN were chosen for managing interaction and rendering within the simulation, as they best suited the requirements of providing a powerful and flexible system and offer many options for functionality.

3.2 Creating the Virtual Environment for Rugby Skills Training (VERST)

3.2.1 The Ball: tracking and modelling solutions discussed

Here, the work directly related to the ball is described. First, the International Rugby Board (IRB) regulations for the shape and design of the ball are examined. With regards to these specifications, a virtual model of the ball is created and textured. The method for tracking with the PST is described: how the optical tracker captures the ball and how this data is imported into XVR and applied to the ball model. Finally, other available potential tracking devices are evaluated.

3.2.2 Creating the Ball Model in XVR

A scaled sphere was used in XVR, to form an ellipsoid that matched quite closely the IRB dimensions (see appendix 7.3). A sphere with a radius of 1 m is created and then scaled to be 30 cm long and 19.74 cm wide along the y - and z -axes; the ball is then textured with the ball texture shown in figure 3.3a, producing the ball shown in figure 3.3b.




Figure 3.3: (a) Ball texture–note texture coordinates are reversed so that it maps correctly in the XVR coordinate space. (b) The ball specifications according to law 2 of the IRB game laws.

3.2.3 The Virtual Reality Peripheral Network

The PS Track software contains a built-in interface to VRPN that allows direct streaming of tracker data to the server. XVR supports VRPN via a set of inbuilt functions (VR Media S.r.l., 2014) that makes it possible to link the PST to XVR via VRPN.

By drawing the ball object at the reported tracker location, it is easy to move the virtual ball synchronised to the real ball. It is possible to set the rotation of the ball using a similar method. There are four data types available from the tracker: axis-angle, angle-axis, Euler and quaternion. Axis-angle was used as the function to set rotation in XVR only takes axis-angle as input.

The starting point for the system was a simple throw of the ball – could a physics system be developed to take the initial flight parameters of the real ball and apply a fully-calculated flight path to the virtual ball?

An experiment was designed to capture a sample of movement data from the PST with the intention of feeding the data into a physics system that processes the real movement and provide a full path calculation. The first task was to see what data

could be captured from the PST; the physics system could then be built around the capabilities of the tracker.

The real ball used was a Skill-tec training ball (Skill-Tec, 2014); this is a size 4 ball often used in training for improving ball handling, and featured a bungee attachment for training exercises (see figure 3.4) attachable to either a ground-swivel or a wrist-strap. Note that a match ball is size 5, but following consultation with Rygbi Innovations Cyf. it was agreed that this size difference would not adversely effect the fidelity of the ball passing training tasks. This suits the requirements well, as a way to restrict the movement of the real ball around sensitive and expensive equipment is essential. Figure 3.5a and figure 3.5b show an experimental setup where the user faces a screen showing the virtual ball. The user holds the real ball which is fixed to the ground or attached to the wrist, and tracked from the side by the PST.

From this experimental setup, it became clear that the data available from the tracker is the position and rotation of the object. Given that the velocity of an object in motion is defined as the rate of change of position (equation 3.1) and that XVR can provide an accurate timer (measured in microseconds), it will be possible to calculate the velocity of the ball.

$$velocity = \frac{\Delta position}{\Delta time} \quad (3.1)$$

3.2.4 The Method for Ball Capture

Figure 3.6 demonstrates the initial design to capture the appropriate data. The PST operates within a capture area noted as the black dashed line. The ball moves across the capture area and through a ‘capture window’ marked here as the red box. As the ball passes through the capture window, it crosses two virtual ‘lines’ (designated ‘A’ and ‘B’) at which time the ball’s data is recorded. This will provide two reference




Figure 3.4: Skill-tec training ball anchored to the ground with a large flat board.

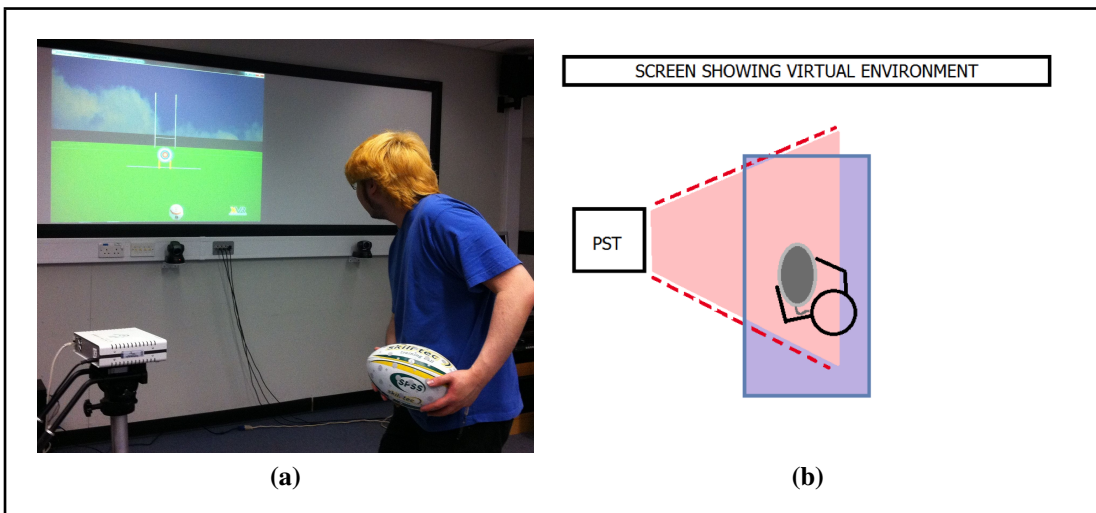


Figure 3.5: (a) Experimental setup of the hardware. (b) Diagram of the experimental setup described in figure 3.5a.

points containing time and position of the ball during crossing that can be used to determine velocity.

If the ball is within the capture area, the motion is tracked. If the x coordinate value

of the ball matches the x coordinate value of the virtual A , the current position and rotation of the ball are recorded along with a timestamp. This group of values – the cartesian position coordinates, the axis-angle rotation and the timestamp recorded as the ball passes a virtual line – will be referred to as the A set and B set. The tracking of the ball continues; if the ball's x coordinate value matches the chosen virtual B and an A set has already been recorded (to prevent reverse-flow), a B set is stored. If both an A set and a B set have been stored, the system calculates a new set that is continuously updated to be the next position. The ball's position and the rotation are updated to the new values on each frame until the ball hits the virtual ground.

Work on the testbed environment began with determining the best location for the A and B virtual lines. By holding the ball at the centre, furthest left and furthest right positions of the capture area, a series of position values were taken from the tracker and used to gain insight on the measuring system of the PST (see table 3.1). The average centre value of -0.009 *PST units* shows the slight error in performing the test with only visual guidance. The left edge of the capture area is -0.22 *PST units* and the right edge of the capture area is 0.23 *PST units*. The same measurements were made




Figure 3.6: The initial design for the capture system.

in real space with the distance being approximately 20 cm in either direction from the centre of the capture area to the edge. It has therefore been taken that the PST's units of measurement are 1 *PST unit* to 1 metre.

Table 3.1: The average data readings for the x coordinate values of segments of the capture area (the left of capture area is physically to the right of the camera which is facing the user, and vice versa for the right of the capture area).

Position	Left of capture area	Centre	Right of capture area
x coordinate values	-0.218347	-0.00920938	0.229973
	-0.218351	-0.00920874	0.229977
	-0.218352	-0.00920809	0.229977
	-0.218341	-0.00920809	0.230014
	-0.218351	-0.00920629	0.230015
	-0.218352	-0.00920629	0.230013
	-0.218354	-0.00920567	0.23001
	-0.218357	-0.00920596	0.230006
	-0.218357	-0.00920502	0.230005
	-0.218358	-0.00920586	0.230002
Average	-0.218352	-0.00920694	0.229999

The ball position could not be directly compared with a single value, however, as it rarely hits an exact value. It was therefore necessary to perform the capture of A and B sets by taking the first value once the ball position had moved past a given point. The code would therefore continually compare the position of the tracked ball to the virtual points – new A sets would be repeatedly stored until the ball's x position value was greater than 0; then it would take the first data received past 0.1 as the B set. After the A and B sets are captured, flags prevent reverse-flow throughout the system.

A third dataset (C set) was then introduced that allowed the determination of an average velocity through three points (two sets of velocity calculations, see figure 3.7).

The 55 Hz capture rate of the PST proved to be acceptable for slow-moving objects, but when attempting to track a ball moving at the speed of a regular rugby pass, the tracker would continuously lose the ball, often not capturing even the three points

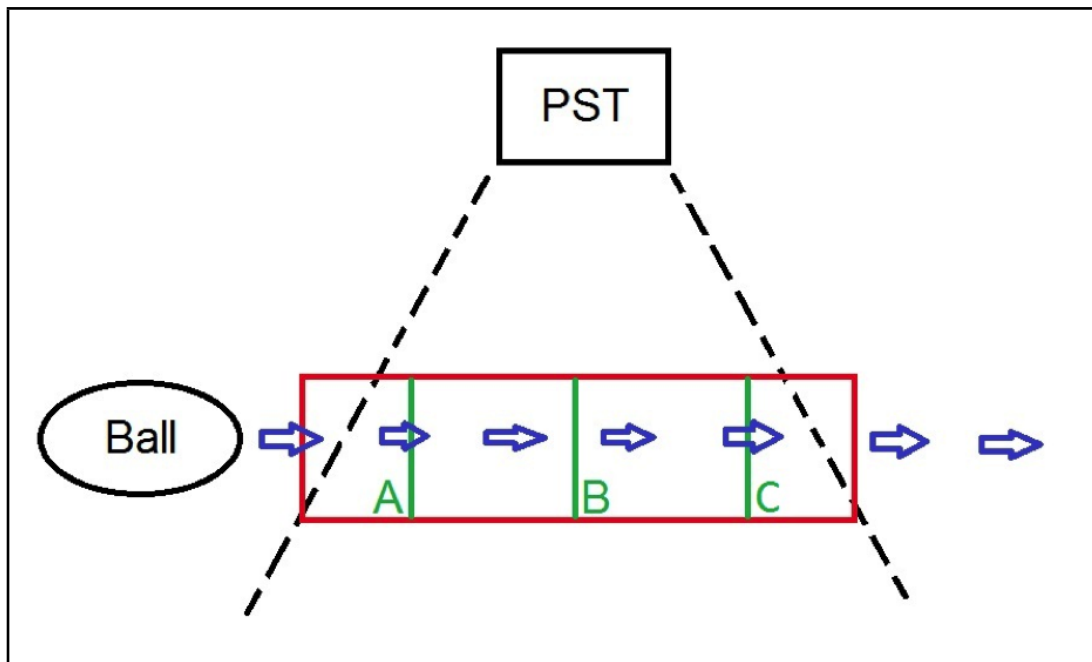


Figure 3.7: New model for data capture, incorporating an extra data set to provide a more accurate representation of the velocity.

needed for the velocity calculations. This problem with the capture rate of the PST prompted experiments with additional types of tracking devices. A Vicon system was available for the experiments, housed at the School of Sport, Health and Exercise Science (SSHES) in Bangor University.

3.2.5 Vicon Motion Capture

The Vicon system was used to investigate the movement of the ball during full passes. A full-sized Gilbert rugby ball was stickered with retroreflective markers (as shown in figure 3.8): three flat markers were placed on each face panel, and a spherical marker was placed on each end of the ball. A simplified version of the data recorded can be seen in the Matlab figure produced in figure 3.9, where the lines are recorded in the same colours as the markers. Here, only the centre markers (marked as circular in figure 3.8) of each panel were plotted. The spin of the ball is visible in the plots – highlighting the importance of spin during a rugby pass – and the ball travels in the expected parabolic arc; the physics simulation will attempt to tackle these motions.

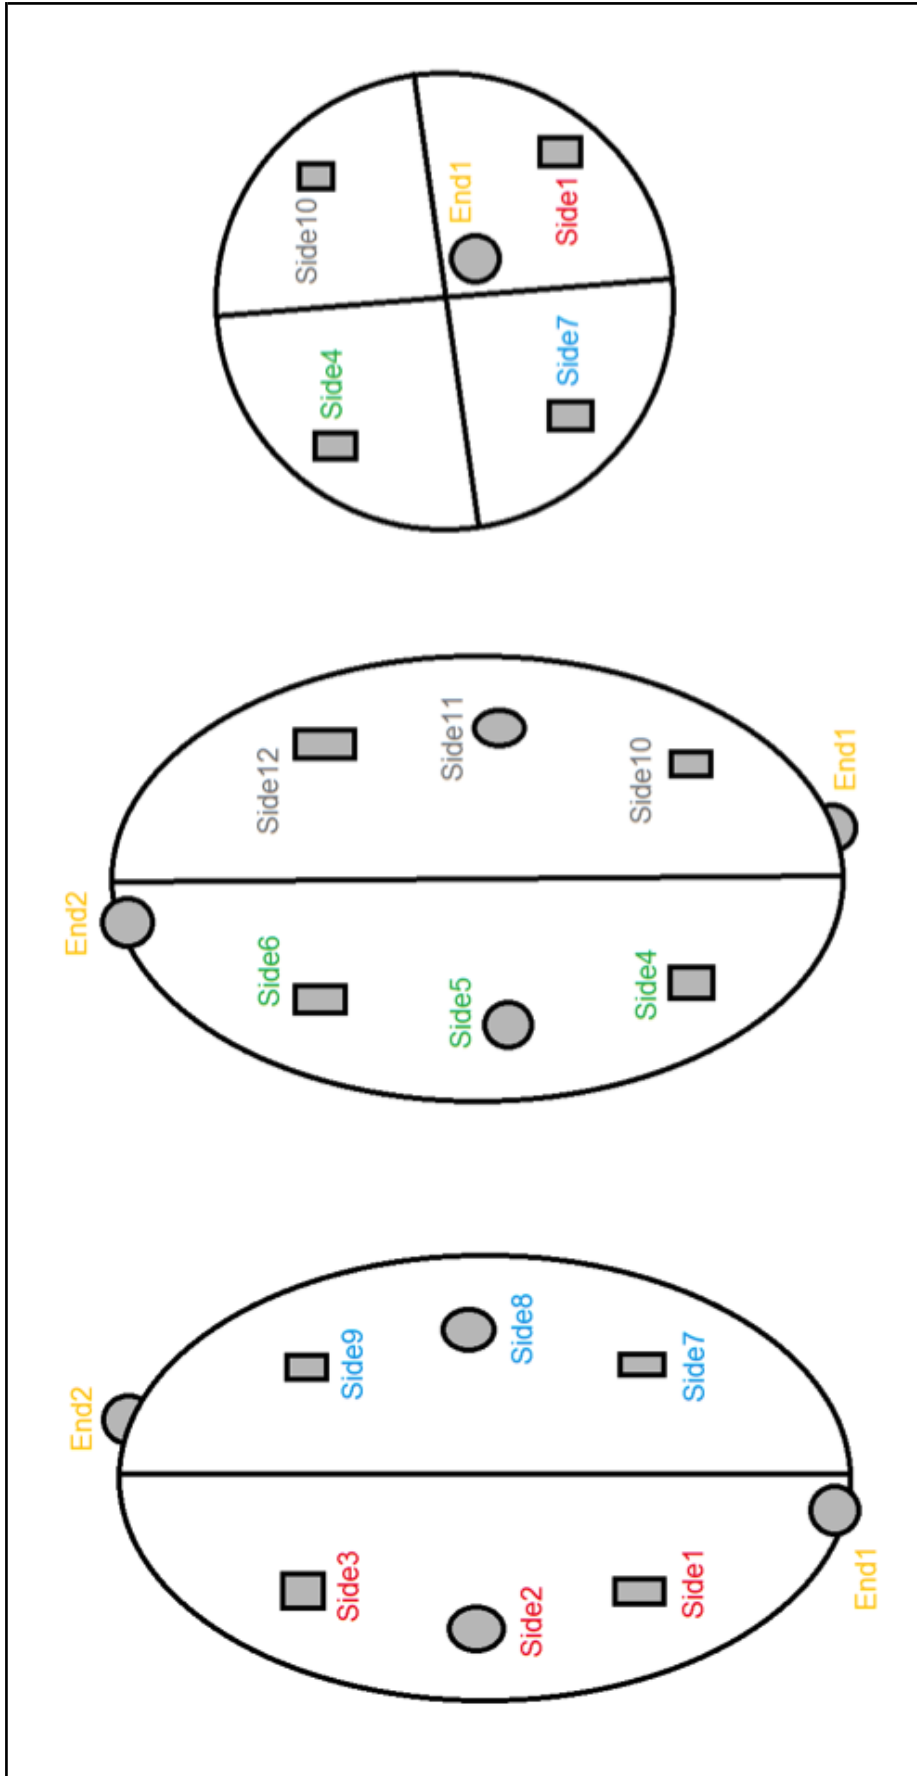


Figure 3.8: Marker-map for the ball during the Vicon capture. The end markers were spherical, side markers were flat, circular markers, and the remaining markers were flat squares cut from a strip of retroreflective tape.




Figure 3.9: Matlab plot of the four panel markers and two end markers.

3.3 Constructing the Virtual Rugby Pitch

First, the setup of the virtual rugby pitch will be described. Details of how the objects are created are given and discussed in relation to the IRB Laws of the Game for rugby union (see appendix 7.2). Further objects are added to the pitch for interactivity: targets, flags and a camera.

3.3.1 The Pitch

Using law 1.2, a basic starting texture was created for use while designing the pitch ground (shown in figure 3.10a).




Figure 3.10: (a) Initial basic texture for the pitch. (b) Posts created with the basic procedural objects available in XVR.

In XVR, a box of size 150 m long, 81 m wide and 10 cm deep was created and the texture was applied. Using the basic cylinder procedural object in XVR, two sets of goal posts were made according to the specified dimensions in law 1.4 (shown in figure 3.10b). First, six cylinders are created: four 15 m tall by 14 cm wide poles for the upright posts, and two 5.6 m long by 14 cm wide crossbar posts. The crossbars are rotated and then each individual post is translated to its correct position on the pitch. Each goalpost set is translated ± 50 m along the x -axis, to place them at each end of the pitch. At each end, the upright posts are each translated ± 2.8 m along the z -axis

to separate them by 5.6 m, and the crossbars are lifted by 2.93 m so that the top of the crossbar is 3 m above the ground.

The textures were updated to look more realistic using free-to-use textures downloaded from www.cgtextures.com. A combination of the textures shown in figure 3.11 were used to create the improved ground texture shown in figure 3.12; the new texture gives the appearance that the pitch surface is not a flat box.




Figure 3.11: Three images used to create the new pitch texture.




Figure 3.12: A new, more complex ground texture created from a combination of the textures in figure 3.11.

3.3.2 The Skybox

To improve the look of the scene, a skybox (sidvind.com, 2014) was used. A skybox is a box or cube that is placed large enough for all of the objects in the scene to fit inside; the inside of the box is textured with images of the sky, providing the illusion that there is a sky behind the scenery. A box was created from an XVR procedural object and made to be $400 \times 400 \times 400 \text{ m}^3$ with the world origin at the centre so that it will surround the user. Figure 3.13b shows the skybox at full size around the pitch and stands. It is now possible to see that the skybox resembles a sky in the background of the scene.

3.3.3 The Targets

The ball passing exercises will involve the player throwing the ball at virtual targets. The first target to be created was a basic archery-style target made up of concentric circles (referred to as the bull's-eye target). A texture was created as seen in figure 3.14a and mapped to a transformed cylinder.

First, the texture is loaded, then a basic cylinder is created with a 50 cm radius and 10 cm height; the caps (circular ends of the cylinder) are selected to appear. Subsets of the cylindrical mesh are then textured: the side of the cylinder is assigned a plain brown material, while the caps are textured with the bull's-eye target. Because cylinders are automatically loaded upright, it is rotated 90° around the z -axis to have the target facing the user.

The second target to be created was the Brian Moore's Mouth target, based on the BBC Wales Scrum V's "Brian Moore's Mouth Challenge" which is an inspiration for the training task that has been selected. The challenge is for the player to pass as many rugby balls as possible (a maximum of 24 balls are supplied) through Brian Moore's mouth in one minute from 10 yards away. The Brian Moore's Mouth target for this




Figure 3.13: (a) The free-to-use sky texture used for the skybox. (b) The scene with the newly-added skybox.




Figure 3.14: (a) Bull's-eye target texture. (b) The source image for the Brian Moore's Mouth target texture. (c) Brian Moore's Mouth target texture.

research was created from figure 3.14b (BBC Wales Scrum V, 2014). The image was cropped and re-aligned to create the texture shown in figure 3.14c.

The texture has been designed with transparent section inside the mouth and above the head to mirror the real target. This was accomplished by creating the transparent areas in the open source GNU Image Manipulation Program (GIMP) and saving the file in the Portable Network Graphics (*.png) format which preserves the transparent areas.

A set of rugby player targets were designed (see figure 3.15) to provide users with a target of which they will be familiar with the dimensions. This target is created in the same way as the Brian Moore's Mouth target. The texture is loaded and used on a box with a $1.818 \times 0.615 \text{ m}^2$ front, and a depth of 0.1 cm. Again, the images were developed in GIMP and saved in the PNG format to preserve transparency.

Each of the final target objects are shown in the scene in figure 3.16a, figure 3.16b and figure 3.16c. Multiple targets in the same scenario are also supported, implemented using an OpenGL Display List.

3.3.4 The Camera

The view of the scene provided to the user was initially designed as a static camera position, but could be extended to allow motion parallax using tracked glasses.




Figure 3.15: (a) Plain rugby player target; (b) blanked-out rugby player shaped target; (c) blanked-out rugby player shaped target holding small bull's-eye target; (d) plain rugby player target holding small bull's-eye target.

The static camera view used for the testbed environment was based on the approximate average height for a Welsh adult male^a, and the targets being placed around the centre of the pitch. The height is 1.8 m (6 ft), and this is added to the y-axis of the pitch position so as to root it to the ground. The camera is set to be 5 m back from the centre line of the pitch (a position of $[-5.0, \text{avg_player_height}, 0.0]$ with a target view of $[50.0, 0.0, 0.0]$). This positions the camera back from the centre of the pitch enough to see the movement of the tracked ball. The target ensures the camera faces the end of the pitch as though the user is looking straight forward (see figure 3.17).

^aActual average measurement 1.770 m (5 ft 9.5 in) of men aged 16+ in Wales, from the Welsh Health Survey 2009 (Welsh Assembly Government, 2009).




Figure 3.16: Each of the different target types present in the scene: (a) the bull's-eye, (b) Brian Moore's Mouth, and (c) the player holding a small target.

3.3.5 The Flag and Windsock

A method of representing any wind effects to the user in the VE was debated within the research group. The most simple method was implemented for the testbed environment: representing wind through the visual cue of a flag or windsock. Although windsocks are not regular features of a rugby pitch, Rygbi Innovations Cyf. decided that they would like it to be explored nevertheless. The other method considered was to use an electrical fan to simulate wind on the user, though it was decided that this would be beyond the scope of this initial prototype system.

To create the flag object, two separate meshes, a flagpole and a flag, were created using cylinder and box procedural objects. The flagpole cylinder is 1.3 m tall and has




Figure 3.17: The static camera view of the pitch.

a 1 cm radius, while the flag itself is a box 33 cm tall and 39 cm wide. It has a depth of 1 mm to reflect the thin material normally used for flags. Once the two meshes are turned into objects, the flag is positioned next to the pole and the pole is added to the flag as a child object to create one larger object under the flag hierarchy.

The windsock was created in accordance with U.S. specifications (Federal Aviation Administration, U.S. Department of Transportation, 2014) to provide realistic measurements. The fabric portion of the windsock is created to be equivalent to a size 2 windsock: 12 feet (3.75 m) long with a 36 inch (0.9 m) throat diameter. The pole of the windsock was made using a 10 m tall cylinder with a 5 cm radius, while the fabric sock was created using a cone frustum 3.75 m tall, with a minimum radius of 30 cm and a maximum radius of 45 cm. The windsock pole is set at ground level, and the fabric section is placed 6 m above ground level which places the top 25 cm lower than the top of the pole. The pivot point of the cone frustum is moved from the centre of the base to the top of the pole and the two parts are kept as separate objects so that the sock can be rotated vertically as well as spinning around the pole. The windsock and flag can be seen in their static positions (no wind present) in figure 3.18a.

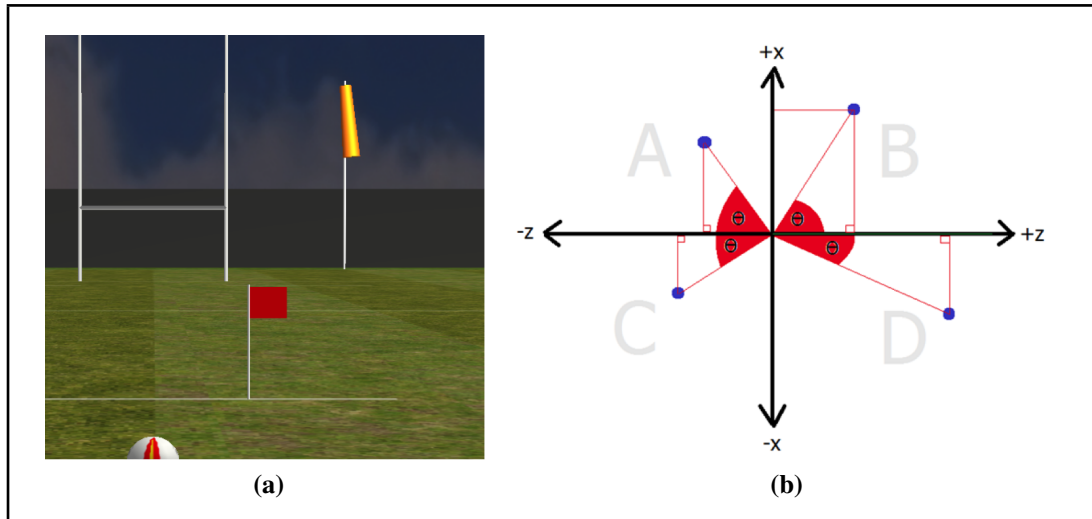


Figure 3.18: (a) The flag and the windsock with no wind in the scene. (b) Illustrates the calculation of the flag direction based on the wind velocity vector.

The wind velocity (set as a vector) is used to rotate the flag and the windsock around their poles, as shown in figure 3.18b (where the 0° mark is along the $+z$ axis, between the B- and D-quadrants). Firstly the wind velocity is categorised into a quadrant, then the angle is returned from the \tan function of the relevant sides and converted from radians to degrees using the pre-defined `RAD_TO_GRAD` constant. The angles are all calculated with the opposite (x) and the adjacent (z) sides of the component triangle: the B-quadrant is simply the angle calculation. The A-quadrant is 180° plus the calculated angle which will be negative because the z must be negative to be in this quadrant. The C-quadrant is also 180° plus the calculated angle, but the double-negative combination of $-x$ and $-z$ results in a positive angle that lies in the $180\text{--}270^\circ$ region. The D-quadrant is 360° plus the negative angle produced from the combination of the $+x$ and $-z$ angles.

A fluttering effect was applied to the flag and windsock because they look unrealistic as static objects that are supposedly being moved by the wind. A loop was created to cause a change in the movement based on the speed of the wind. A counter variable is increased every time the draw loop is run (every frame). When this frame count reaches a specific number (initiated as zero but decided by the wind velocity) a random

number in a specified range is taken and run through a check to change the direction of the flutter (make the value positive or negative) each time it is called. At the end of the loop, the counter variable is reset to zero.

Once the wind velocity magnitude has been calculated, the windsock's vertical movement is set. The rotation is grouped into ranges of wind speed: 0–1.5 ms⁻¹, 1.5–3.5 ms⁻¹, 3.5–5.6 ms⁻¹, 5.6–7.6 ms⁻¹ and 7.6+ ms⁻¹. These ranges are based on the U.S. specifications (Federal Aviation Administration, U.S. Department of Transportation, 2014) which state that the windsock should be perpendicular to the pole for wind over 15 knots, 28 kmh⁻¹ or 17 mph and should “move freely about the vertical shaft it is attached to and when subjected to wind of 3 knots (5.6 kmh⁻¹ or 3.5 mph)”. Because the flag and windsock are not moving across the screen, the velocities are suitable in metres per second (that is to say, they do not need to be converted to metres per frame). The three velocity measurements provided for each requirement (15 knots, 28 kmh⁻¹ or 17 mph, and 3 knots, 5.6 kmh⁻¹ or 3.5 mph) each convert to different amounts in metres per second, so the smallest value is taken for each: 7.6 ms⁻¹ for the perpendicular movement and 1.5 ms⁻¹ for free movement. Extra ranges were added between these values to provide a range of vertical rotations for the windsock.

For each range, the windsock is rotated around the pole then vertically (around the *x*-axis), the previous rotations are cleared, then the rotate function is used to add the *x*-axis rotation on top of the *y*-axis rotation. The flutter speed is then set because it is the number of frames to leave between random movements, the lower the number the more frequent the flutter movement, therefore higher flutter rates are used for higher wind speeds. The variables holding the amount of the flag and windsock random movement for the current frame are also updated: the random numbers are calculated between zero and the movement variables, so the higher the value the larger the movements. For the flag the higher wind speeds cause smaller but more frequent flutters; for the windsock higher winds cause larger and more frequent vertical movements.

3.4 Additions to the Environment

Now that the basic environment had been set up and populated with objects, the next step was to develop the functionality to control the behaviour of the scene and the objects within it.

3.4.1 Interactivity

3.4.1.1 Hitting the Virtual Targets

When including targets, it is necessary to check if the ball has actually hit a target on throwing.

It began as a simple case of checking whether the target and ball shared the same space at any frame of the throw. The current position of the target was retrieved and then used to create a bounding box around the target based on the scale as shown in figure 3.19. The position of the ball was then tested to see if it fell within the bounding box of the target; if true the target was considered to be hit.




Figure 3.19: The target position being used to create a bounding box around the target.

A problem emerged that, in several trials the ball passed through the target but failed to activate the target being hit. This was because the ball was meeting the bounding box conditions in the yz -plane but the ball was being drawn either side of the target without intercepting it (see figure 3.20).




Figure 3.20: Here, the conditions for intersecting the target bounding box were met on yz -plane (left) but not the xy -plane (right)

To counteract this problem, the bounding box was extended along the x -axis as shown by the blue box in figure 3.20 and the previous ball position and the current ball position were tested relative to the target.

Extending the bounding box of the targets along the x -axis now allows the target hits to be correctly registered. It is hoped that the complexity of the conditions will ensure there are no false-positives recorded where the ball falls short of the target.

3.4.1.2 Moving Targets

In order to provide a potential range of training scenarios, a method of applying simple movement to any target in the scene was created. The basic movements were organised into a 'movement schema' shown in table 3.2 and figure 3.21.

The target movement requires three variables: the reference number for the movement type, the velocity of the target and the range of movement the target will have (the

Table 3.2: The movement schemas available for targets.

Reference Number	Direction of Movement	Colour in figure 3.21
0	Static (no movement)	-
1	Left–Right	Green
2	Up–Down	Orange
3	Forwards–Backwards	Purple




Figure 3.21: The movement schemas in table 3.2, (a) shows the target from the front, (b) from the side.

distance span the target will be able to cover). Each of these variables are contained in vectors to support multiple targets.

3.4.1.3 Ball Trails

The ball trails function was created as an aid to visualising throws (see figure 3.22). They are created using a series of small spheres in a display list. The ball trail point object is created during scene set up as a 3 mm red sphere. Once the ball hits the ground, a new trail sphere is placed at each position. The resulting trails appear as in figure 3.22.




Figure 3.22: Ball trails being demonstrated with a tracked throw.

3.4.2 Administrative Controls

3.4.2.1 System Reset

A system reset function was devised so that it was not necessary for users to leave the program to make another pass attempt. The function simply resets specific variables in the program when it is summoned, and this is enough to set the system back to the starting point.

3.4.2.2 Pre-prepared Throw (Debug Mode)

In Debug Mode, the tracking code is bypassed and a dummy ball vector is used for demonstration and administrator testing. The ball is set up with an initial position of $[0, 0, 0]$ and an initial velocity of $[5, 5, 0]$, which produces a nice parabolic curve from the physics engine (see figure 3.23).

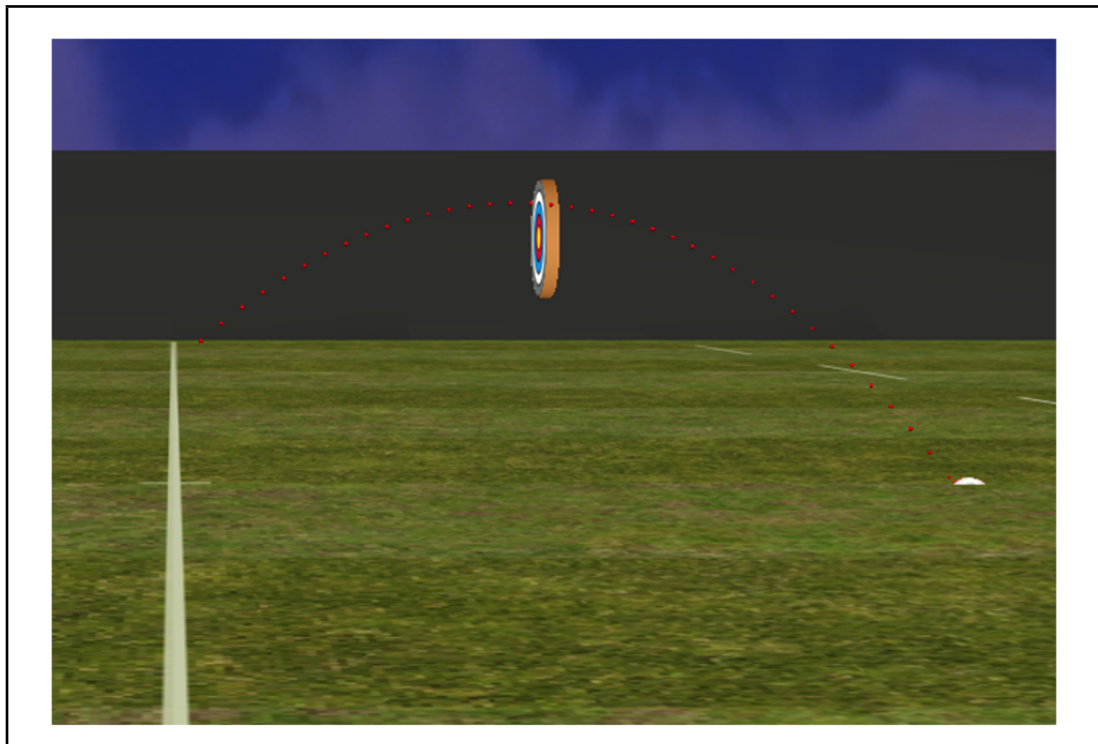


Figure 3.23: The debug throw shown using ball trails.

3.4.3 User Interface

3.4.3.1 Keyboard and Mouse Control

Keyboard and mouse controls were added to VERST to provide access to both user and administrator options to customize the environment. The keyboard shortcuts are summarised in table 3.3. It should be noted that the Num Lock on the keyboard must be activated for the Num Pad keys to work, and that the Debug Mode key-press of ‘d’ has been programmed to be insensitive to case.

Table 3.3: Key-presses versus the function or task that they trigger.

Key Assigned	Function/Task
Space	Reset the whole system for another throw
Num Pad 0	Static main user camera view
Num Pad 2	Camera: facing centre of pitch from the side.
Num Pad 4	Camera: facing opposite direction from target from centre of the pitch.
Num Pad 5	Camera: top-down view of the pitch.
Num Pad 6	Camera: target view.
Num Pad 7	Camera: left goal close-up.
Num Pad 8	Camera: facing centre of pitch from the side.
Num Pad 9	Camera: right goal close-up.
F1	Cycles through the type of target used (bull’s-eye, Brian Moore or rugby player).
F2	Toggles the ball trails on/off.
F3	Toggles the Settings console in the top-right corner of the screen on/off.
D	Enters Debug Mode – runs a pre-programmed throw.

The VERST menu (figure 3.24a) is an important aspect of the user interface, as program users cannot be expected to remember keyboard shortcuts. It is also a more visually accessible way of providing options. Right-click menus are regularly featured in programs, making it a logical way to provide options to the user.

If the settings are changed, they can be instantly reset to the default option with the “Default Display Settings” option. The system is put in Tracker Mode with the static user camera view; trails, the settings console and the flag camera are disabled. The ball is drawn shaded and textured, and is shown pre-throw. Finally, the system is reset and prepared for a new throw.

3.4.3.2 Statistics Console

The statistics console is a small text box programmed to appear once a throw has been completed to provide the user with some basic statistics about their last pass (see figure 3.24b).



Figure 3.24: (a) The options menu displayed when the right mouse button is clicked on-screen. (b) The statistics console. (c) The settings console.

The total time is taken at the end (the time between the timer being activated when the *A set* is taken and the moment the ball hits the ground), the magnitude of the total

distance is calculated and an average velocity is calculated from the two. The data is stored and compiled into useable strings ready for drawing to the scene.

3.4.3.3 Display Options Console

A small settings window can be drawn in the top-right corner of the screen if required (see figure 3.24c). The window can be toggled between being active or inactive either by key-press or selecting an option in the click-menu.

Once the window has been drawn, text is added to reflect the current options selected by the user. After the “Settings” title which remains unchanged, the first line details the current mode of the system: tracker or debug. In a similar way, the status of the ball trails option is recorded: if trails are active then “Trails ON” is printed, but trails are off by default. Finally, on the lowest line of the settings console, the current camera angle is noted. The camera angle number recorded refers to the Num Pad key of the camera.

3.5 Rugby Ball Physics Engine

The challenge of the physics engine was to produce the most accurate system that could run successfully in real time to serve the graphical interface with the movement data. While it would be possible to eventually derive a system that considered all factors, the system would require a large amount of processing power: too much for a standard PC to run in real-time, without even considering the processing load caused by the graphical portion of the program. It was necessary to sacrifice certain aspects of the system and to make some assumptions to reach a compromise between smooth, speedy processing and functional fidelity.

The physical motion of the ball is being modelled on the following aerodynamic principles: the trajectory is a straight line unless external forces intervene; the initial velocity depends on the launch angle; on ascent, weight and drag act together, on de-

scent the forces have opposite signs; any external side forces like wind or rain are added to the lift. In practice, the ball's flight is a complex motion of a solid object through a constantly moving but slow paced fluid. Any changes in the physical parameters of the fluid (viscosity, density, internal pressure) due to weather conditions change the resistance force that acts on the ball during flight. Hence, in order to increase the simulations accuracy, the air is modelled as a Newtonian fluid with variable viscosity depending on velocity and temperature. The interaction between solid and fluid is time dependent. During flight, the ball rotation is very complex (three rotation axes similar to an aeroplane), depending on the initial throw conditions and the motion stability in its entirety.

Several iterations of the physics engine are described in this section, starting with the initial straight-line prototype to test the ability to process information from the tracker. The method for calculating the velocity of the ball is explained, and the progression of the physics engine to a parabolic curve. In accordance with the section describing the development of the tracking process, the addition of a third set of data points (the *C set*) and the required code adaptation is discussed. The first complex model is explained with reference to the background literature used in its development, followed by details of how the advancement of this model could include some basic weather effects in the future.

3.5.1 The Straight-Line Prototype

The physical system began with the implementation of a simple straight-line movement that confirmed the possibility of using the *A* and *B sets* to feed into a loop that could calculate the next movement step for the ball.

The first equation used was the classic slope-intercept form linear equation (Weinstein, 2014a):

$$y = mx + c \tag{3.2}$$

where m is the gradient (slope) and c is a constant. The m is calculated as

$$m = \frac{\Delta y}{\Delta x} \quad (3.3)$$

With the A and B sets, it was possible to calculate the gradient and constant to determine the rest of the flight path. After the A and B sets are taken, the new ball position is calculated. The starting position of the ball is set to the last point it was tracked at, then two delta-values are calculated for the x and y coordinate values. A delta-value is not calculated for z , as sideways movement was not considered at this stage.

The two delta-values are then used to calculate the gradient. The equation is reversed to find the constant using the A set and the newly derived gradient. The difference in time between the A set and the B set being taken, and the distance travelled is also calculated to give the speed of the ball (although the speed is not used in this version as it has been calculated in distance per second and the draw cycle is 100 frames per second). For the physics module to run, both A and B sets must have been taken, and the x coordinate value is incremented at a steady speed of 0.0001 units/frame (approximately 0.01 ms^{-1}). The y coordinate value is then calculated using the equation and the position of the ball object is set.

Once the A set and B set have been taken, the ball travels along a straight line path from the release point (as soon as B is taken) at the gradient it began with between the capture points. The physics code loops 50 times to ensure it does not run infinitely.

3.5.2 Velocity Calculations

It was necessary to develop a way to accurately measure the velocity of the ball, and to correctly display the velocity on screen based on the frame rate of the application. As discussed earlier in the chapter (on page 65), the velocity is calculated as the delta distance divided by time taken (equation 3.1).

But because the XVR timer reports the time in microseconds (μs) the times produced must be converted to seconds by dividing the timers by 1000000 ($1s = 1 \times 10^{-06} \mu s$). Unfortunately, this leads to a loss of precision in XVR, resulting in timers being rounded to whole numbers. Given that the throw is over in less than one second, a high precision is essential for the velocity calculations. To counter this problem, the distance can instead be multiplied by 1000000; this maintains an accurate result. Then the number of frames needs to be considered – there are 100 frames per second in the draw loop, so the velocity needs to be divided into distance/frame.

$$\text{total number of frames} = \Delta time(s) \times 100 \quad (3.4)$$

This can be used to derive the virtual velocity expressed in distance units per frame:

$$\text{velocity (units/frame)} = \frac{(B_{[x]} - A_{[x]})}{\left(\frac{(B_{time} - A_{time})}{1000000} \times 100\right)} \quad (3.5)$$

$$\frac{(B_{[x]} - A_{[x]}) \times 1000000}{((B_{time} - A_{time}) \times 100)} \rightarrow \frac{(B_{[x]} - A_{[x]}) \times 10000}{(B_{time} - A_{time})} \quad (3.6)$$

3.5.3 The First Basic Curve

A basic parabolic curve was programmed to be the next step of the flight-path development (see figure 3.25). L represents maximum length of the flight and H is the maximum height. The equation for the parabolic flight path is defined as follows, beginning with a standard quadratic equation:

$$y = ax^2 + bx + c \quad (3.7)$$

where the height of the curve (y) is defined as the squared length (x) multiplied by some constant a added to the length multiplied by a constant b plus a constant c . At the




Figure 3.25: Planned parabolic flight path of the ball.

start of the curve, both x and y are zero, meaning that the constant c must also be zero. At the end of the curve, x is now the maximum length L , but y is zero again. The rest of the equation becomes $aL^2 + bL = 0$. If the whole equation is divided by L , it can be reduced to $aL + b = 0$.

$$\left\{ \begin{array}{l} x = 0 : y = 0 \\ x = L : y = 0 \end{array} \right. \rightarrow \begin{array}{l} c = 0 \\ aL^2 + bL = 0 \end{array} \rightarrow aL + b = 0 \quad (3.8)$$

Rearranged, it becomes the following:

$$b = -La. \quad (3.9)$$

Replacing the b in the original equation results in this next equation:

$$y = ax^2 - aLx \rightarrow y = ax(x - L) \quad (3.10)$$

Now that the length has been worked into the equation, the maximum height must be too. At the point where the y is the maximum height H , the x is half the maximum length.

$$x = \frac{L}{2}, \quad y = H \quad (3.11)$$

By replacing the x and y in equation 3.8 with those from equation 3.11, equation 3.8 becomes:

$$H = a \frac{L}{2} \left(-\frac{L}{2} \right) \rightarrow H = -a \frac{L^2}{4} \rightarrow a = \frac{-4H}{L^2} \quad (3.12)$$

The final equation is expressed below. It should be noted that the sideways (z -axis) movement is still not yet being considered:

$$y = \frac{4H}{L^2} \times (L - x) \quad (3.13)$$

The equation was successfully deployed in XVR, producing a neat parabolic curve based upon the velocity and direction of the captured data.

3.5.4 Adding The *C Set*

As a result of the issues with the capture rate of the PST, a third set was added to provide the ability to take an average velocity and provide a more accurate ball velocity. The previous set size was halved, allowing a new set (named the *C set*) to be added (see figure 3.7). The *A set* is still the last point captured before 0.0 *PST units*, the *B set* is now captured between 0.0–0.2 *PST units*, and the *C set* after 0.2 *PST units* (the previous *B set* capture location).

3.5.5 Rugby Ball Dynamics

Bray and Kerwin (2003) worked on the development of a mathematical model for the flight of a soccer ball, in the context of a direct free kick; figure 3.26 shows the accompanying diagrams.

In figure 3.26, the curved dashed line depicts the flight path of a soccer ball kicked along the y -axis. Figure 3.26 (a) shows the ball leaving the ground at origin O . At point P , the ball is at an angle of Ψ from the (x, y) plane and travelling with a velocity of v (in a direction described by unit vector $\boldsymbol{\tau}$). In figure 3.26 (b), the ball (still at point P) is at an angle of θ away from the y -axis, and in figure 3.26 (c) the spin of the ball is demonstrated. It is assumed that the ball spins at a constant angle, denoted γ , an inclined axis parallel to the (x, z) plane (the ball is pitching). Unit vector $\boldsymbol{\sigma}$ specifies the direction of the γ -spin.

The simplest form of the motion, considering only the forces, can be written as follows:

$$\mathbf{F} = m\mathbf{g} + \mathbf{F}_d + \mathbf{F}_l \quad (3.14)$$

The resultant force \mathbf{F} acting on the ball can be determined by the mass (m), multiplied by gravity (\mathbf{g}) and adding the drag (\mathbf{F}_d) and lift (\mathbf{F}_l) forces. The drag and lift forces are defined as follows where ρ is the air density, A is the cross-sectional area, and C_d, C_l are the drag and lift coefficients:

$$\mathbf{F}_d = -\frac{1}{2}\rho Av^2 C_d \boldsymbol{\tau}, \quad (3.15)$$

$$\mathbf{F}_l = \frac{1}{2}\rho Av^2 C_l \boldsymbol{\sigma} \times \boldsymbol{\tau}. \quad (3.16)$$

v is the magnitude of the vector \mathbf{v} , calculated as the square-rooted sum of the squared Cartesian components:




Figure 3.26: Diagrams by Bray and Kerwin (2003).

$$v = |\mathbf{v}| = \sqrt{v_x^2 + v_y^2 + v_z^2} \quad (3.17)$$

These previous equations are combined to produce the differential equation for the flight of the soccer ball (\mathbf{r} is the position, $\dot{\mathbf{r}}$ is the velocity and $\ddot{\mathbf{r}}$ is the acceleration):

$$\ddot{\mathbf{r}} = \mathbf{g} - k_d v^2 \boldsymbol{\tau} + k_l v^2 \boldsymbol{\sigma} \times \boldsymbol{\tau} \quad (3.18)$$

Where:

$$\ddot{\mathbf{r}} = \frac{d^2 \mathbf{r}}{dt^2}, \quad k_d = \frac{\rho A C_d}{2m}, \quad k_l = \frac{\rho A C_l}{2m} \quad (3.19)$$

$$\boldsymbol{\tau} = \cos \Psi \sin \theta i + \cos \gamma \cos \theta j + \sin \Psi k \quad (3.20)$$

$$\boldsymbol{\sigma} = \cos \gamma i - \sin \gamma k \quad (3.21)$$

with i , j and k as unit vectors on (x, y, z) . Three equations were derived to each solve the \ddot{x} , \ddot{y} and \ddot{z} components of the acceleration based on the velocity magnitude and components $(\dot{x}, \dot{y}, \dot{z})$:

$$\ddot{x} = -v\{k_d \dot{x} - k_l \sin \gamma \dot{y}\} \quad (3.22)$$

$$\ddot{y} = -v\{k_d \dot{y} + k_l [\cos \gamma \dot{z} + \sin \gamma \dot{x}]\} \quad (3.23)$$

$$\ddot{z} = -\mathbf{g} - v\{k_d \dot{z} - k_l \cos \gamma \dot{y}\} \quad (3.24)$$

Since these equations rely on the unknown first derivation of the components, it is not possible to solve them. They can, however, be approximated using an Runge-Kutta 4th Order (RK4) method given an initial set of conditions. Since the C_d and C_l parameters for a soccer ball had only been estimated in previous work, a C_d of 0.25 and C_l of 0.23 was determined by Bray and Kerwin through an experiment; it is assumed that the C_d and C_l remain constant.

While this work provided a starting point for the equations of motion of a ball, the parameters would need to be adjusted to cater for a rugby ball rather than a soccer

ball. The aerodynamics of a rugby ball have been researched by Vance et al. (2012). In a wind tunnel experiment, a series of values for the lift and drag coefficients of the ball were measured. The drag and lift coefficients were recorded for yaw angles from 0–60°, and were found to be functions of the Reynolds number. They state their results are consistent with the experimental results obtained by Seo (2004) and Alam et al. (2009).

Here, a problem is encountered in that there are varying potential values for the drag and lift coefficients, and it is not clear which should be used for the model being designed for this research. An assumption is made that the yaw movement (Ψ) of the ball will be very small in a pass. For yaw angles of 10–20° the drag and lift coefficients have very similar ranges of values. While they are shown to vary with the Reynolds numbers, it is by a very small amount (< 0.05). This lead to the values of $C_d = 0.25$ and $C_l = 0.23$ being maintained.

3.5.6 Model Development

Once the tracker velocities have been set the physics model continuously recalculates the position of the ball until the ball reaches the height of the ground plane.

The 6-dimensional vector *ball* holds three position components and three velocity components in the following format:

$$ball = [pos_x \quad vel_x \quad pos_y \quad vel_y \quad pos_z \quad vel_z] \quad (3.25)$$

The y coordinate value of ball is compared with the y-component of the pitch, to see if the ball is still above ground. If so, the physics model is run. The physics model repeats a simple set of instructions:

1. pass the current values to the `Integrate()` function (see appendix 7.4.1),
2. increment the time by the given time-step dt ,

3. set the ball object's position from the new values,
4. clear the buffers,
5. draw the ball at the new position,
6. redraw the scene.

The `Integrate()` function performs an approximation to the results of equations 3.22, 3.23 and 3.24 using the RK4 method.

The RK4 method is performed as follows (Weisstein, 2014c): four increments ($\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4$) are made by estimating a slope function and a weighted average taken to approximate the next step ($y_{(n+1)}$):

$$y_{n+1} = y_n + \frac{1}{6}(\mathbf{k}_1 + 2\mathbf{k}_2 + 2\mathbf{k}_3 + \mathbf{k}_4) \quad (3.26)$$

$$\mathbf{k}_1 = hf(t_n, y_n) \quad (3.27)$$

$$\mathbf{k}_2 = hf(t_n + \frac{1}{2}h, y_n + \frac{1}{2}\mathbf{k}_1) \quad (3.28)$$

$$\mathbf{k}_3 = hf(t_n + \frac{1}{2}h, y_n + \frac{1}{2}\mathbf{k}_2) \quad (3.29)$$

$$\mathbf{k}_4 = hf(t_n + h, y_n + \mathbf{k}_3) \quad (3.30)$$

For each of the four incremental steps, a time (t_n) and state (y_n for \mathbf{k}_1 , $y_n + 1/2\mathbf{k}_1$ for \mathbf{k}_2 , $y_n + 1/2\mathbf{k}_2$ for \mathbf{k}_3 , and $y_n + \mathbf{k}_3$ for \mathbf{k}_4) are passed to the function, the returned value is multiplied by the step-size (h). The `Integrate()` function takes the current state of `ball` and the step size `dt`. The RK4 method is applied to all six individual elements of the `ball` vector where the position elements are converted into velocity and

the velocity elements are converted into acceleration. For *k2*, *k3* and *k4*, the previous step (*k1* for *k2*; *k2* for *k3*; *k3* for *k4*) is used for the calculation. The starting point for each of these steps is taken to be the result from the previous step by taking each of the six `state` elements and adding it to its derived element multiplied by the time step; that is, the position elements are:

$$\mathbf{position}_{xyz} = \mathbf{position}_{xyz} + (\mathbf{velocity}_{xyz} \times \Delta time) \quad (3.31)$$

and the velocity elements become:

$$\mathbf{velocity}_{xyz} = \mathbf{velocity}_{xyz} + (\mathbf{acceleration}_{xyz} \times \Delta time). \quad (3.32)$$

After *k4*, the process of adding the time-step-multiplied integral is applied to the original ball using the *k4* values. These new position and velocity values are the new *state* of ball, and once it has been determined, the *ball* vector is returned to the physics model to update the position of the ball.

The four *k*-increments are found by sending the current *state* of ball to a function which performs the approximation: `Rigid1()` for *k1*, and `Rigid2()` for *k2*, *k3* and *k4*.

For *k1*, the `Rigid1()` function (which can be seen in appendix 7.4.2) extracts the current velocity (lines 4–6) and finds the magnitude (for use in the equations of motion, line 9), then the *output_dball1* vector (the delta position and velocity, lines 13–20) is filled with the results.

For *k2*, the `Rigid2()` function (which can be seen in appendix 7.4.3) receives the current state along with the time-step value and the resulting *k1* vector. A vector named *new_ball* is calculated as each element of the current state added to the corresponding *k1* element multiplied by the time-step (lines 4–16). It then performs the same calculations as the `Rigid1()` function (lines 14–30). The `Rigid2()` function repeats this process for *k3* and *k4*, using the relevant previous increment vector in the

calculations. It should be noted that the formulae for y and z have been swapped to fit the coordinate system of XVR (in which the y -axis is vertical).

3.5.7 Adapting the Model for Weather Data

Adding weather effects to the system could be achieved by modifying the drag and lift coefficients used in the calculation of the flight path of the ball. Rain would dampen the surface of the ball and lower the drag coefficient by making it slicker, while wind would affect the drag and lift of the air, depending on the direction and strength of the wind. The coefficients would need to be measured using experiments in the field to provide realistic results.

Similarly, models of pitches at different locations around the world could be modified to include air pressure effects on the flight path of the ball. Temperature effects on the ball could also alter the surface properties.

Although the use of weather data was not accomplished during this research, it could be added to the system in the future.

3.5.8 Testing the Model using Vicon Data

To test the results of the physics engine, a series of real throws were captured in a motion capture laboratory at Bangor University. The laboratory was equipped with 12 Vicon cameras covering an area of approximately $5 \times 5 \text{ m}^2$ in the centre of the room, allowing only space for relatively short throws. The cameras were able to capture a full throw with sufficient accuracy and precision, at a rate of 250 Hz. Once the data had been collected, the throw was analysed in a Matlab implementation of the physics engine, the results of one of the throws can be seen in figure 3.27. The red line represents the real captured data and the blue lines the predictions of the physics engine given the real data prior to the start of each line. In the top-down view, the small fluctuation between 2–3 m is caused by the ball spinning.




Figure 3.27: Matlab implementation of the physics engine in action. The red line represents the data of the real tracked ball; the blue lines each represent a predicted flight path by the physics engine given the real data up to the point from where they start.

The physics engine required a sample of points to calculate an average velocity. Reading the points from the beginning onwards, the physics engine took the first 20 points initially to calculate line (a) in figure 3.27, then 40 points to calculate line (b), incrementing by 20 each time until 120 points were used for line (f). Figure 3.27 demonstrates the substantial difference in distance that can occur: the real ball landed at approximately 6.5 m, between trajectories (c) and (d), but only trajectory (d) follows the real trajectory quite closely; trajectories (a), (b) and (c) were too short, and (e) and (f) were too long. The recorded motion does not include data on when the ball left the hands of the user and entered free flight, which would be the ideal time to capture the motion of the ball. The curve upwards before the beginning of (b) shows the ball still in the hands of the user, and not yet at the free flight velocity required for the most accurate physics engine results. The result of (f) reveals the simplicity of the physics engine, in which it has averaged the entire set of real data into a single velocity vector and therefore has an incorrect snapshot of the throw. Figure 3.27 also demonstrates the effect of capture errors on the trajectory, particularly the top-down view in the lower graph.

This testing of the physics engine demonstrated the necessity of controlling the capture window to ensure the correct section of the real ball's flight is captured. Different methods of restricting the point of data capture were considered; while there was the 'capture window' method described earlier in the chapter (on page 66), the use of conductive fabric on the surface of the ball, the use of small Arduino electronics systems and cyber gloves were discussed, but it was decided they would be too restrictive on the user. This led to the design of the ball capturing system to be a small capture window placed approximately 50 cm in front of the user. With the participant in a fixed position, this would allow only the capture of the ball after being released from the participant's hands. Restricting the size of the window to 1-1.5 m would concentrate the captured data to the initial portion of the flight where the ball has the initial free-flight velocity.

3.6 Summary

In this chapter, VERST has been extensively designed and features implemented for a testbed system created using the XVR software package, although implementations using other software environments is also possible.

The testbed featured several aspects not previously used in other research: a method of using a real rugby ball as a haptic device in the VE; a method for capturing position, orientation and time data from the moving ball; and a new adaptation of Bray and Kerwin's physics model to calculate the flight of the rugby ball.

The physics engine for the rugby ball in flight employed a simple ballistic model based on work by Bray and Kerwin (2003) on football free kicks, and by Vance et al. (2012) on the aerodynamics of a rugby ball in flight. Using the mathematical model proposed by Bray and Kerwin and drag and lift coefficients based on the work of Vance et al., the solution to the flight path of the rugby ball can be estimated in real-time using the RK4 method. The required input data for the physics engine to function was the last position vector at which the ball was recorded for continuity between the real and computed movement, and a Cartesian velocity vector which was taken from an average of the several velocity estimates taken as the ball was tracked. The data produced by the physics model was found to be accurate providing the ball movement is captured at the exact moment the ball was released from the users' hands.

The following design requirements were given prior to the development of VERST:

- 3D content authoring to create the virtual scene;
- stereoscopic rendering to explore whether the use of stereo improves the VE;
- free to use; to allow for uses with a limited budget;
- ability to develop on a Windows 7 64-bit PC (the PC provided for the development work);

- software that can accept a variety of input devices to allow the upgrading of the tracking hardware;
- provide tools to produce a program which can be run easily by non-specialists;
- the program produced should not be resource intensive and, if possible, run platform-independently, since the specifications of the future running system are unknown.

VERST has successfully fulfilled the design requirements. A virtual scene that can be viewed in stereo has been created for VERST with the free-to-use authoring tool XVR, which can be run on Windows 7. It is possible to use a large variety of input devices through the VRPN interface. VERST, while currently lacking an extensive user interface, can be started simply by opening the executable file. VERST is not resource intensive, and has been successfully used on PCs with lower specifications than the one on which it was developed.

Some of the features designed and implemented in the testbed system were not carried forward for use during experimentation. The system was kept at a very simple level for experimentation, and features deemed to be unnecessary for basic experiments were the different target types (only the Bull's-eye was used); the flag and windsock, as wind was not used in the physics engine; moving targets; balls trails; the administration menu and options; and the statistical panel.

To investigate the suitability of this environment, an experiment would be required. It would be necessary to study the different components designed in this chapter, as well as the ability of people to use the system. The next chapter describes the first experiment carried out.

Chapter 4

Experiment 1: An Initial Evaluation of the Efficacy of VERST

A testbed system had been developed to attempt to determine the answer to the hypothesis “A virtual environment can be used to effectively train ball passing skills in rugby” With the completion of the testbed system, the attempt to answer some of the research questions described in chapter 2 could begin.

An exploratory experiment was designed to provide an initial insight into the capabilities of the system whilst also attempting to investigate the system in the context of the following research questions:

- What types of skills appear to be best suited to training in VEs?
- Which technologies achieve the best results in specific performance measures?
- Should stereoscopic displays be used?
- What does it cost?

Thanks to the Visionair (2014) project, the VERST research activities at Bangor University was partnered with a research group in Aix-Marseille Université for use of their CAVE, but because of the difference in the facilities it was necessary to port the

system. Visionair holds partnerships with a number of universities and research centres all over Europe. It seeks to pair these partners with groups who submit applications for access to use the facilities. The group working on this project submitted an application to Visionair, and was accepted for collaboration with the Centre de Réalité Virtuelle de la Méditerranée (CRVM) at the Université de la Méditerranée: Aix-Marseille II in Marseille, France. The pairing was appropriate, as CRVM's previous projects are a series of virtual environments, some of them sport-based; the centre has strong links to the Institute of Movement Sciences at the university.

This section describes the challenges in porting the whole system to different software and hardware and what experiments were possible with the system running with higher performance capabilities.

4.1 Description of CRVM Facilities

CRVM houses a CAVE environment within the centre. The CAVE consists of four projected screens (the floor and three walls), each wall screen is 3 m wide and 4 m tall, and the floor is 3 m squared. The two side screens (henceforth referred to as the *left* and *right* wall screens) are moveable, which allows the CAVE to be configured in several ways (see figure 4.1 (a)). Figure 4.1 (b) shows the fully closed CAVE from above; figure 4.1 (c) shows the CAVE with both walls opened, the grey areas represent non-projected floor space. In figure 4.1 (d), only one of the walls is open (the L-shape configuration), and figure 4.1 (e) shows one of the walls opened to 45°. The three wall panels must remain attached along the 4 m edges, otherwise the two separate walls can use any combination of fully opened, fully closed and half-way to create the desired space.

Each screen had a resolution of 1400×1050 pixels and the screens were seamlessly joined to provide a visually continuous projection surface. For this experiment, the L-




Figure 4.1: The CAVE configuration possibilities.

shape configuration was used to allow the user more space and for efficient placement of the motion tracking system.

Stereoscopic projection is achieved by two DLP projectors per screen. Stereoscopic separation between left and right eye images is ensured by colorimetric separation (an Infitec technological solution). Infitec filters are installed in the projectors, while the subject wears glasses with the same filters. The system uses wavelength multiplexing to separate projected colours into pairs of wavelengths (one per eye) and provide an effective passive stereoscopic presentation.

A tracking system (ART), using infrared recognition of passive markers placed on the subjects glasses, was used to record the subject's head position and orientation and to update the stereoscopic images relative to the subjects point of view in real time (with a frame rate of 60 Hz). The two eyes' images were presented simultaneously and the images were correctly geometrically projected for each eye. The whole projection system is controlled by a cluster of six high-performance synchronized PCs.

A real rugby ball was suspended from the ceiling of the building (see figure 4.3a), tracked by five Vicon cameras (see figure 4.3b) indicated in figure 4.2 with 'V' markers. The participants were positioned at either *A* or *B* (marked in figure 4.2); the white squares in figure 4.2 represent two platforms that were raised approximately 10 cm above the ground along with the floor screen of the CAVE. The capture area (marked

by the red hatched box in figure 4.2) was positioned far enough away from the user that they would have to reach unnaturally to have their hands within it. This ensured that data would be captured from the ball after the participant had released it.




Figure 4.2: A top-down view of the screen configuration. The grey sections highlight the screens, whereas the ‘A’ and ‘B’ markers show the two positions the participant stands during the experiment. The ‘V’ markers give the position of the Vicon cameras and the red hatched area indicates the approximate capture area.

4.2 Experimental Design

The experiment was originally designed in consultation with Dr Simon Watt due to his background in experimenting with distance perception and stereoscopy in VEs. Initially, the experiment was designed to investigate three key factors and their effect on perceived distance^a:

^aThe three factors were taken directly from draft of experimental design by Dr Simon Watt.




Figure 4.3: (a) The tethered ball suspended from the ceiling with a rope. (b) One of the Vicon cameras used in the experiment.

1. stereoscopic versus non-stereoscopic presentation;
2. the fidelity of monocular cues (such as ground-plane texture); and
3. the position of the physical display, whose surface structure is specified by focus cues, and a visible outline.

These factors were investigated by modifying the following conditions within the VE:

1. stereoscopic presentation;
2. a floor screen; and
3. user's physical distance from the screen.

Distances for the VE were originally as follows: the user would be positioned at three distinct distances from the screen – 1.5, 3 and 6 m – to provide a good comparison

for the effect of physical distance, and the distance from the user to the target would be between 5–15 m.

Several different methods of conducting the experiment were considered, based on previous work in the field (Interrante et al., 2006). Verbal estimation (Napieralski et al., 2011) of the target distances could be taken from the participants: a simple experiment to perform, but subjects may experience difficulty quantifying their motor responses, and their response may not accurately reflect the estimates made by their motor system. Walking to the perceived distance is a common distance perception experiment (Thomson, 1983; Willemsen et al., 2009; Steinicke et al., 2009), where the subject walks blindly to a previously-seen target. While this estimate would come directly from the subject's motor response, it would be difficult to perform with a fixed screen – previous blind-walking experiments in a VE used HMDs (Interrante et al., 2006; Willemsen et al., 2004; Creem-Regehr et al., 2005), which allow the subject to move up to the virtual position without space restrictions. Distance matching was also considered, where a real target is moved to match the distance the user perceives the virtual target to be at (Williams et al., 2008); while this is an intuitive process for the subject, external cues – such as height in the visual field – may cause distortion in the estimates. The chosen method of conducting the experiment was to simply engage the system as intended: the subject would throw the ball to hit the target (Sahm et al., 2005). While this would require the physics engine to be estimating the flight path of the ball as accurately as possible, it would also be useful to see the complete system working as intended.

Four participants were recruited for this study, all rugby playing students at Aix-Marseille Université, but hardware restrictions would not allow the subjects to be placed at multiple distances: the Vicon cameras set to capture the flight of the ball operated in a small volume (approximately $1.5 \times 1.5 \times 1.5 \text{ m}^2$) and the participants could not move closer to or further from the screen. However, it was possible to position the participant either side of the outer floor screen corner (as shown in figure 4.2).

Based on this information, the experiment was re-designed. The key factors are the same first and third as originally considered:

1. Stereoscopic versus non-stereoscopic presentation;
2. The position of the physical display, whose surface structure is specified by focus cues, and a visible outline.

The original second factor was abandoned to concentrate the experiment on the suitability of the hardware. The experiment conditions were updated to reflect the removal of participant distance from the screen:

- Position of the participant relative to the screen;
- Stereo/non-stereo;
- Floor projection on/off.

The experiment was run with the system configured in a combination of these conditions, with each participant doing 35 trials for each configuration (7 different target distances, with 5 repetitions of each at different rotations around the participant's azimuth), as shown in table 4.1. Each participant performed the configurations in a different order to guard against bias in the results (see appendix 7.5.1 for full specification).

4.3 Developing the Experiment

While the basic Virtools version of the VERST system had been developed beforehand, the system was enhanced and the experimental controls were not developed until the procedure had been fully designed.

The Vicon and ART tracking systems were both utilised. The Vicon system tracked the real ball while the ART tracked both the user's glasses and provided secondary

Table 4.1: Trial Configuration table for experiment 1.

Configuration Number	Participant Position (A/B)	Stereo / Non-Stereo	Floor / No Floor
1	A	S	F
2	A	S	NF
3	A	NS	F
4	A	NS	NF
5	B	S	F
6	B	S	NF
7	B	NS	F
8	B	NS	NF

tracking for the ball. The rugby ball and glasses with the rigs holding the retroreflective markers used for tracking the two objects are shown in figure 4.4a and figure 4.4b.

While the glasses were already equipped with the rig, the rugby ball was prepared specifically for the VERST system. A double-ended screw was fixed into the point at the end of the ball where the four edge seams meet. The rig was then screwed onto the screw and fastened with a nut. The other end of the ball was secured to the ceiling with a rope, as can be seen in figure 4.4a.

The ART system was not capable of capturing the ball alone, as the ball's rig faced away from the ART cameras the majority of the time, so the Vicon system was vital for capturing the movement of the ball accurately (see figure 4.4c). The Vicon system, however, was not able to capture the glasses for the same reason, requiring the use of both tracking systems in tandem. This led to an issue where the Vicon cameras would occasionally mistake the glasses for the ball as a throw was being made and the participant obscured the ball from the cameras behind them while bringing the ball upwards and moving their head downwards. This was solved by covering the rear of the markers on the glasses with small pieces of tape, hiding them from the Vicon camera while keeping them visible to the ART cameras in front of them.

A guidelines and consent form for the participants was prepared explaining the




Figure 4.4: (a) The tracking rig suspending 5 retroreflective markers at the end of the rugby ball. (b) Infitec glasses with retroreflective markers on the periphery of the lenses. (c) The semi-opened CAVE with the ball suspended from the ceiling of the building and five Vicon cameras surrounding the ball.

procedure of the experiment, and questioning participants on their experience playing rugby, which hand is dominant, if they have ever had any eyesight problems, whether they wear glasses/corrective lenses, or are colour-blind. It also asked whether they were happy to be filmed during the experiment. As the participants were all French, the guidelines and consent form were translated into French^b to ensure full understanding and the participants asked questions and corresponded with the collaborating researcher running the experiment, with responses being translated back into English.

A full test-run of the experiment was conducted with one of the collaborating technicians as the subject. The results were analysed and an issue with the processing of the tracked data was found where the average velocity was becoming skewed by erroneous processing of the times from the Vicon system. This problem was corrected and the system was found to be functioning acceptably.

4.3.1 Porting of the System to Virtools

All of the CAVE equipment had been configured to work with the Virtools software environment. It was straightforward to port the VERST software to Virtools and we decided to do this rather than calibrate all of the hardware with XVR. The Millennium Stadium model Model created by user 'dizzyHARSH' (2014) (figure 5.2a) used in the CRVM experiment was downloaded from the Google Sketchup Warehouse and imported into Virtools.

To ensure certain measurements were consistent between the XVR and Virtools environments (for experiments), it was decided a configuration file should be developed. Each portion of the conversion will now be described in detail.

^bTranslation work done by Vincent Perrot, a researcher at the CRVM facility.

4.3.2 Developing the Configuration File

The final configuration file can be seen in figure 4.5. Lines 1-22 are comments explaining each of the variables. The `playArea` has been set as 150 m long, 1 cm deep (below ground), and 81 m wide, with the `playOrigin` being at (0,0,0). The main camera view is 0.8 m above the origin. The pitch mesh ground level is at -1 m making the camera 1.8 m above the ground, which is approximately 6 ft, chosen to reflect the average height for a Welsh adult male (Welsh Assembly Government, 2009). It is also pulled back 5 m to ensure the ball is visible whilst being tracked by the PST. Three targets are initialised by the `no_targets` variable. The `target_scale` is set to be 1 for each of the three targets. The positions of the targets are as follows: target 1 is 10 m away, 1 m above the origin (2 m above the ground) and directly in the centre; target 2 is 5 m away, 0 m above the origin (1 m above the ground) and also directly in the centre; target 3 is 20 m away, 5 m above the origin (6 m above the ground) and 2 m to the right of the origin.

The example file tests a variety of movement schemas: target 1 remains static while target 2 operates on a left-right movement and target 3 up-down. The target velocities are measured in metres per frame, meaning that target 2 moves at 1 cm/frame (approximately 1 ms^{-1}) and target 3 moves at 3 cm/frame (approximately 3 ms^{-1}). The range of target 2 allows it to move left-right ± 2.5 m either side of the original position defined in `target_pos`, and target 3 moves ± 1.5 m above and below the original position.

The configuration file is loaded by the software renderer (XVR or Virtools). If not found, then a set of default values will be used.

4.3.3 Adaptation of the Physics Engine

A UDP connection can be used to send and receive information from and to the Virtools environment. Using this mechanism, the physics code was implemented as a

```

1  /*
2
3  VERST Configuration File
4  -----
5
6  playArea    define total area of play (size of pitch)
7
8  playOrigin  set the point of origin in the virtual area of play
9  camOrigin   set the default camera location within the pitch
10
11 noTargets   number of targets on the pitch
12 targetScale set the scale of the target (where 1 is normal-sized
13             radius of 0.5m)
14 targetPos   assign target position values
15 targetMov   specify a target movement plan: 0 static; 1 sideways;
16             2 height; 3 forwards/backwards
17 targetVel   set the velocity of the target
18 targetRange total distance the target will cover (will split
19             either side of the starting point)
20
21 windVel     velocity of the wind
22 flagPos     position of the flag on the pitch
23
24 NOTE: make sure there are enough target attribute entries to
25       correspond with the number of targets
26 */
27
28 playArea = [ 150 0.01 81 ];
29
30 playOrigin = [ 0 0 0 ];
31 camOrigin = [ -5.0 0.8 0.0 ];
32
33 noTargets = 3 ;
34 targetScale[noTargets] = { 1 1 1 };
35 targetPos[noTargets*3] = { 10.0 1.0 0 5.0 0.0 0.0 20.0 5.0 -2.0 };
36 targetMov[noTargets] = { 0 1 2 };
37 targetVel[noTargets] = { 0 0.01 0.03 };
38 targetRange[noTargets] = { 0 5 3 };
39
40 windVel = [ 4 0 2 ]
41 flagPos = [ 70.0 -1.0 35.0 ]

```

Figure 4.5: The configuration file with example values. Line numbers are shown for reference.

standalone module. Once Virtools has received the coordinates of the tracked ball from the Vicon system, it processes the coordinates into an array set equivalent to the ball vector (equation 3.25) previously discussed in chapter 3 (on page 98).

Once this has been done, Virtools passes the data over UDP to the standalone VERST physics engine. The physics engine processes the ball array and predicts a position for the ball at every time step until the ball hits the ground. This data is stored in an array in the following format and returned to Virtools over UDP.

$$result = pos_{x1}, pos_{y1}, pos_{z1}, \dots, pos_{xn}, pos_{yn}, pos_{zn} \quad (4.1)$$

The physics engine is implemented using the same method as the XVR prototype described in chapter 3.

4.4 The Experiment

4.4.1 Experimental Procedure

Before the participant arrived for the experiment, the system was prepared. The projectors were turned on and required approximately 10 minutes to warm up. The ART and Vicon systems were individually calibrated with wands then both calibrated with a single floor square to match the coordinate systems (shown in figure 4.6a). The two tracking systems' synchronisation was verified – and sometimes corrected – in Virtools. Finally the tracking of the ball and glasses were tested; the ball's tether was checked and a few test-throws were briefly made to test the system and verify regular operation.

For each participant the procedure was the same: when they arrived at the facility they were welcomed and given the information sheet (see appendix 7.5.2) explaining the experiment. Once they had read through the information they were given the opportunity to ask any questions and were asked to fill in the participant's questionnaire



Figure 4.6: (a) Calibrating the ART system with the calibration wand. (b) The Essilor Digital CRP used to measure participant's interpupillary distance.

(see appendix 7.5.2). They were tested for stereoacuity in the normal range using the Randot SO-002 (Stereo Optical Company, Inc., 2014). The Randot SO-002 provides a series of geometric forms which the participant must identify from the random dot stereogram backgrounds that they are contained within to ensure the stereoacuity of the participant falls between 20–500 arcseconds (0.3–8.3 arcminutes).

Following the stereoacuity test, the participant's Interpupillary Distance (IPD) was measured using the Essilor Digital Corneal Reflection Pupillometer (CRP) (Essilor, 2014) (shown in figure 4.6b). The IPD was then used to tailor the stereo presentation to the subject by displaying the two individual images at the given separation and ensure the correct geometrical projection for each eye.

The participant was then provided with a rugby ball of the same construction as the one used within the system – a Skill-tec size 4 training ball – and asked to pass the ball between themselves and one of the technicians for 5 minutes. The purpose of this real-world exercise was partly to give them a chance to get used to the size of the ball (smaller than the regular size 5 ball used by adults) and partly to allow the observation of their passing skills. This practice session was conducted indoors to avoid interference from the wind – which did not feature in the experiment – and space restrictions only allowed for a passing distance of 9 m. However this was considered

sufficient, as it was well-centred within the range of targets used in the experiment (5–15 m).

The participant was then shown into the CAVE area and was allowed time to experience throwing the tethered ball without the system activated. Once they felt prepared, the lights were turned off to maximise the participant's attention to the screens, and all non-essential personnel vacated the testing room. One colleague was permanently positioned at the control station (see figure 4.7a) with a view of the system (see figure 4.7b) to manage the experiment. The experiment could be paused at any time if the participant was not prepared for the trial (the ball's tracking rig often became tangled in the tether).




Figure 4.7: (a) The control station for the experiment, situated next to the CAVE, a technician would be present at the station during the entire experiment. (b) The view of the VERST system from the control station. The CAVE walls are marked by the grey checkered planes. The trial details are displayed in the top-left corner of the window.

At the start of the experiment the participant was positioned at the relevant position to their initial configuration. The trials were ordered so that all configurations in a single position were performed in one go. Once the first four configurations had been completed, the system would pause to allow the participant time to change position, and would only recommence when the technician commanded.

Once the experiment began, the program would automatically assign the seven target distances randomly between 5 and 15 m at random angles rotated around the

position of the participant (obtained from the tracked glasses) and each configuration would be cued one after another. A beep would signify the start of the trial, and the participant would have 5 seconds to attempt the throw. This ensured that the participant could not have time to study the target and contemplate the distance and to encourage the participant to react to their first instincts. If the user failed to throw the ball in time or the Vicon cameras did not successfully capture the motion of the ball, they were awarded one further attempt at the same trial. Despite there being an existing physics model, the virtual ball was not shown to the user so that they could not make any corrections to their method or adjust their attempts to compensate for any failings in the system. The participant was prevented from knowing whether they had successfully hit any of the targets during or after the experiment. After the participant had completed the experiment, they were asked to provide some verbal feedback on the system.

4.4.2 Participants

Four participants agreed to take part in the experiment. The details of the four participants are tabulated in table 4.2, followed by details and observations recorded about each participant during the experiment.

Participant 1 had taken part in previous (unrelated) experiments in the CAVE and was also present during parts of the development of the VERST system. His familiarity with the hardware was apparent as he was more relaxed than the other participants. During the real-world passing task, it was noted that his passing was very accurate with a fast spin used to keep the ball in a straight line. The participant's comfort in using the system meant that he was happy to throw the ball very hard, although because his throw was considerably stronger than the colleague who performed the test run, it was unforeseen that the ball would so often become entangled in not only the tether, but also the rope suspending the tether to the ceiling. He would also require several seconds to retrieve the ball, as it would swing back far due to the force with

Table 4.2: Details of the participants recruited for the experiment.

Participant ID	Age	Gender	Years playing rugby	Experience (position of play)	Left/right-handed	Wore glasses/contact lenses	Other eye problems	Colour-blind	Stereoacuity in normal range	Consented to being filmed	Date of participation
1	25	M	15	Flanker	R	N	N	Y	Y	Y	25/03/13
2	38	M	32	Scrum half	R	Glasses for Astigmatism; not worn during play	N	N	Y	Y	26/03/13
3	26	M	13	Wing (junior team); currently scrum half in Fédéral 3 (fifth division)	R	N	N	N	Y	Y	26/03/13
4	21	M	10	Inside centre	R	N	Had an operation for myopia 3 months ago	N	Y	Y	27/03/13

which he was throwing it. As the experiment went on, the participant became visibly frustrated by the tangling and the effort required in catching the ball before the next trial. One of the markers on the rig was pulled off as it tangled with the rope, but no damage was caused. When the participant was asked to give feedback, he said the system was ok, admitting that the ball getting tangled caused a lot of trouble. He also

claimed that he had tried to adapt to a gap in the Vicon's capture window (in which he was himself partially masking the ball) but that it did not work so he went back to his original method.

Participant 2 was visibly the most experienced player. During the real-world task, his pass was noted to spin down the central axis of the ball (see figure 4.8a) rather than the skewed one (figure 4.8b) that caused the ball to wobble slightly (which the other three participants did). Participant 2's throws caught in the rope far less than participant 1 and did not suffer from the same wild swinging motion on return. It was noted that his throws would spin on the small metal joint connecting the ball to the tether rather than spinning on the rope and whipping it in circles. However, he had to pause for a moment halfway through the experiment to allow the tether to release the rotation stored in it when the metal joint could not spin fast enough to match the ball spin. Feedback from participant 2 was that he found it strange to not see the virtual ball moving on the screen. He also claimed that he had not noticed the configurations changing as he was focussed on the targets. Despite this, he noted that sometimes targets appeared closer than he thought they should have when the stereo was not active.




Figure 4.8: An axis of spin down the centre of the ball (a) against an axis of spin at approximately 25° causing a skewed spin (b).

Participant 3 was observed with a slightly off-centred rotation angle, but with slower spin than participant 2. Participant 3's throws were very forceful, to the point where the tracking rig detached from the ball twice during two of the trials, once hit-

ting one of the wall screens, and the other time striking the stand of one of the Vicon cameras behind the participant. Participant 3 noted as feedback that he had used pitch lines in the virtual scene as a guide to distances and that he had noticed no difference between either the stereo or floor being on or off.

Finally, participant 4 performed the real-world passing task with confidence and a strong throw down a slightly off-centred angle of rotation. However, during the experiment, participant 4 threw very softly in comparison to the other participants. Since he had been observed with the same level of ability in throwing long distances as the other participants, there was clearly an issue of confidence in using the system. He appeared to be throwing to just beyond the screen distance regardless of the distance of the target, an observation supported by his real ball velocity statistics (mean velocity [3.95, 1.68, -0.16], standard deviation [0.59, 0.56, 1.27]). During the experiment he seemed to grow in confidence marginally and throw slightly harder until the Vicon failed to capture a throw, at which time he would revert to the slower throws. During participant 4's experiment, there were loud noises around the outside of the building as a vehicle moved outside, which may have been a distraction to the participant. Feedback provided by participant 4 was that he had difficulty with the tethering rope catching on his shoulder as he drew back to make the throw. He thought the distances all seemed "ok" and noted that while he noticed the floor screen being turned on and off, he had not noticed the change in stereo on/off.

Participant 1's comment about attempting to adapt to the system highlights one of the problems this system must deal with: the user should not adapt to any issues as they will be practising the skill incorrectly. Since none of the other participants did the same thing, it is possible that participant 1's previous experience with the system may have prompted him to attempt to 'help' by 'making the system work' rather than allowing the system to do the work correctly itself. Participant 2 found it strange not seeing the virtual ball, which was a concern as it gave the VE a surreal quality. It is also possible that seeing the ball fly upwards after being thrown (being pulled up by

the tether) may have been detrimental to the attempt to keep the VE appearing close to reality. Participants 2 and 3 both alleged that they had not noticed the change in stereo and floor conditions, while participant 4 claimed not to have noticed the change in floor usage. This is of interest – if their performance alters with the configurations they may be unconsciously reacting to the depths differently under different conditions using their motor systems, if their performance does not alter then the conditions truly did not appear different to them. Participant 2, despite claiming to not have noticed the difference in conditions, thought that target distances seemed compressed without stereo. This may have been a reaction to viewing the non-stereo scene after his two initial configurations being shown in stereo. That he believed the targets were appearing closer than they should implied that he was using another cue to determine the distance of the targets, which may have been other visual cues in the scene such as the size of the target or the lines on the pitch, although this can only be speculated. Participant 3 stated that he had used visual cues in the scene – the pitch lines – to guide him to the correct distances, which is undesired as he may have been considering the guided distance rather than his motor system’s assessment of the distance. This observation suggests he did not have confidence in the targets visually appearing to be at the correct distance, and was searching for assistance in verification. The system layout with the ball tethered to the ceiling caused problems for participants 1 and 4, but not 2 and 3.

4.4.3 Data Gathered

The data were gathered as a single text file for each participant’s configuration group. A sample of the trials captured for participant 1 is shown in figure 4.9 (with values rounded from 6 decimal places to 1 decimal place for readability).

As shown in the sample, the dataset gathered contains the participant ID number, the date and time the experiment started and the configuration number of the trial group in the header. Then for each trial: the trial number; target position; head position and orientation (taken from the tracked glasses); the final ball position and the average ve-

Participant ID: 1 Date: 03/25/2013 10:15 Trial Configuration: 1					
Trial No.	Target Position	Head Position	Head Orientation	PhysEng Position	PhysEng Velocity
0	10.4, 1, 3.1	0.2, 1.6, -0.6	0.7, 1.2, 0.4	1.5, 1.4, 0.1	7.0, 2.2, 2.0
1	13.1, 1, 4.0	-0.3, 1.6, -0.7	0.7, -0.3, 0.1	-1, -1, -1	-1, -1, -1
1	13.1, 1, 4.0	0.2, 1.6, -0.5	0.9, 1.1, 0.6	1.4, 1.3, 0.1	7.1, 2.6, 2.2
2	...				
8	13.7, 1, -0.6	-0.2, 1.6, -0.6	0.8, 0.5, 0.5	0, 0, 0	0, 0, 0
8	13.7, 1, -0.6	0.1, 1.5, -0.6	1.5, 1.2, 1.2	-1, -1, -1	-1, -1, -1
8	13.7, 1, -0.6	0.1, 1.5, -0.5	1.5, 1.2, 1.3	1.5, 1.3, -0.4	8.4, 3.0, 0.1
	...				

Figure 4.9: Example output file from experiment 1.

locity sent to the physics engine. In trials where the ball was not successfully captured, the physics engine data is recorded as 0, 0, 0. In trials where the participant ran out of time to make an attempt, the physics engine data is recorded as -1, -1, -1. For each successful trial, a separate text file was generated containing the ball's position as recorded by the Vicon tracking system, accompanied by the velocity calculations made between each step. An example of the file contents for participant 1's trial 0 is contained in figure 4.10.

0.58197, 1.06721, -0.16509	000ms	7.09216, 2.69814, 1.92072	032ms
0.808919, 1.15355, -0.103627	064ms	7.16238, 2.37874, 1.83223	016ms
0.923517, 1.19161, -0.0743113	080ms	7.08644, 2.21939, 1.97081	016ms
1.0369, 1.22712, -0.0427784	128ms	7.02156, 2.0278, 1.89824	048ms
1.26159, 1.29201, 0.0179654	144ms	7.04, 1.78313, 1.98368	016ms
1.37423, 1.32054, 0.0497043	160ms	6.77, 1.82687, 2.18843	016ms
1.48255, 1.34977, 0.0847191	176ms	0, 0, 0	000ms

Figure 4.10: Example trial data from experiment 1.

In this file, the first column contains the coordinates taken at the timestamp in the

second column. The third column holds the velocity calculated between the point in the same line followed by the point in the following line, followed by the time difference between the two points in column 4. The final two entries in the third and fourth columns do not have a second point to compare with, thus they are empty. The average velocity as shown in the previous file is calculated from these velocity measurements. The greater the number of points collected, the more accurate the representation of the velocity, provided the capture window is not so large that the ball is captured beginning downward flight distorting the velocity provided by the user (discussed in chapter 3, page 102). The mean number of points collected in the capture window was 9 with a standard deviation of 3 (minimum of 2, maximum of 24), providing accurate data for the physics engine. Each participant was also filmed and/or photographed taking part in the experiment.

Tabulating (table 4.3, see also figure 4.11) the failed trials demonstrated that there was a substantial amount of failed trials that required a reattempt: of the 1120 trials to be performed, an extra 318 (28.4%) had to be performed, with 213 (19.0%) being due to the system failing to capture the ball properly. Many of the timed out trials were as a result of the participant needing time to get hold of the swinging ball, meaning that the majority of the failed trials can be attributed to system failures. There were a larger number of failed captures from configurations 1–4 (150) than 5–8 (63), this is most likely due to the participant often partially masking the ball from the Vicon cameras when in position A.

4.5 Results

The actual data obtained was used by a routine (written in Matlab) that calculated the trajectory of the ball using the physics engine and generated a plot to visualise the throw (shown in figure 4.12a, figure 4.12b and figure 4.12c). In the plot, the green square represents the floor screen; the small red circle is the location of participant 1;

Table 4.3: Number of failed trials during the experiment which required a reattempt.

Participant		Configuration								
		1	2	3	4	5	6	7	8	Total
1	Timed Out	7	8	4	2	10	16	1	0	48
	Failed Capture	8	7	8	4	10	9	5	1	52
2	Timed Out	3	5	5	5	0	3	2	4	27
	Failed Capture	5	12	23	18	4	3	1	0	66
3	Timed Out	4	2	2	3	1	4	1	1	18
	Failed Capture	6	2	0	15	2	3	3	5	36
4	Timed Out	0	0	1	0	3	2	1	5	12
	Failed Capture	9	8	13	12	2	2	3	10	59
Total	Total	42	44	56	59	32	42	17	26	31
	Timed Out	14	15	12	10	14	25	5	10	10
	Failed Capture	28	29	44	49	18	17	12	16	213




Figure 4.11: The failed trials from table 4.3 plotted as a stacked bar chart; this highlights the large number of capture failures compared to timed out trials.

the large red circle shows the position of the target; the trajectory of the ball is marked in two parts, with the real captured ball in blue and the predicted flight path in red.

Once the Matlab routine had plotted the flight path of the ball, it provided the final position of the ball on the ground (the initial hit to the ground, rebounding was not considered). Following these calculations, data derived from original readings were:

- target-participant distance;
- the distance travelled by the ball;
- the distance ball landed from the target;
- the trial result;
- participant-screen distance; and
- the target rotation.

A visual assessment of the result of the trials was recorded from each of the Matlab plots, designating the trial as one of the following:

- **Hit** – the ball’s flight path passed through the target;
- **On Target** – the throw was in the correct direction but did not travel the correct distance (either too short or too far);
- **Miss Left** – the ball was thrown with too much rotation along the azimuth to the left (no detail was recorded about the amount of over-rotation or whether the ball was thrown the appropriate distance);
- **Miss Right** – the ball was thrown with too much rotation along the azimuth to the right;
- **Failed** – no data was recorded in the trial.




Figure 4.12: (a) Participant 1, configuration 1, trial 1 visualised in Matlab. The large red circle is the target; the small blue circle is participant 1's head position; the green square shows the location of the floor screen; the purple section of the flight path marks the captured real data while the yellow portion marks the resulting flight path calculated by the physics engine. (b) The y - z plane of (a). (c) The x - z plane of (a).

4.5.1 Observations and Ideas Explored

Initially there were a number of connections between variables to be explored; table 4.4 holds each pair of variables to be compared followed by basic correlation testing between the two and a comment about the results. The entire dataset was submitted for correlation testing, and while some of the paired variables in table 4.4 had been sought, others were unexpected results.

The first connection to be tested was table 4.4 (6); a strong connection between the mean speed of the tracked ball and the distance the ball travelled as calculated by the physics engine supports the results of the physics engine. The flight of the ball will often take approximately (or just less than) one second to reach the ground thrown from a standing position, and so the speed in metres per second provides a very quick estimation of the travelled distance.

In table 4.4 (1) and (2), a connection between the configuration and the distance thrown or mean speed of the ball was sought, but none was found; suggesting that the configurations may not have affected the distance thrown at all. Further inspection into (9) position, (14) stereo and (3) floor screen found that there was no correlation between any of these factors and the distance the participant threw. A moderate correlation was found between the participant and the distance thrown (8).

A high correlation between the position and the number of misses left (10) and right (11) were found, and the number of misses left and right (7) were moderately correlated. The number of misses left and right had no significant correlation with stereo (15, 16) or floor (4, 5), showing position to be the only influence.

The connection between screen distance, the distance thrown (12) and mean speed (13) were investigated but not found to be correlated, suggesting that the participants were not being influenced by the accommodation to the screen distance.

Finally the target distance and the distance the ball landed from the target were found to be highly correlated on a positive scale, suggesting that the further the target, the less distance the ball travelled.

Table 4.4: Results of correlation testing between variables 1 and 2 (sorted alphabetically by variable 1).

ID	Variable 1	Variable 2	Correlation	Comment
1	Configuration	Distance Thrown	0.004	No significant correlation was found.
2	Configuration	Mean Speed	-0.06	No significant correlation was found.
3	Floor Screen	Distance Thrown	0.05 (F), -0.05 (NF)	No significant correlation was found.
4	Floor Screen	# Misses Left	0.06 (F), -0.06 (NF)	No significant correlation was found.
5	Floor Screen	# Misses Right	-0.06 (F), 0.06(NF)	No significant correlation was found.
6	Mean Speed	Distance Thrown	0.83	As expected, the mean speed from the tracked data is highly correlated to the distance thrown calculated by the physics engine, supporting the results of the physics engine.
7	# Misses Left	# Misses Right	-0.51	Moderate correlation.
8	Participant	Distance Thrown	-0.56	Moderate correlation.
9	Position	Distance Thrown	-0.13 (A), 0.13 (B)	No significant correlation was found.
10	Position	# Misses Left	-0.76 (A), 0.76 (B)	High correlation suggests that the number of misses to the left is related to the position the throw was made from (B).
11	Position	# Misses Right	0.57 (A), -0.57 (B)	Moderately high correlation suggests that the number of misses to the right is related to the position the throw was made from (A).
12	Screen Distance	Distance Thrown	0.02	No significant correlation was found.
13	Screen Distance	Mean Speed	-0.02	No significant correlation was found.
14	Stereoscopy	Distance Thrown	0.06 (S), -0.06 (NS)	No significant correlation was found.
15	Stereoscopy	# Misses Left	0.04 (S), -0.04 (NS)	No significant correlation was found.
16	Stereoscopy	# Misses Right	-0.04 (S), 0.04 (NS)	No significant correlation was found.
17	Target Distance	Distance Ball Landed from Target	0.77	High correlation; suggesting that the further the target, the less distance the ball travelled.

4.5.2 Statistical Analysis

The straight distance thrown by all four participants, as a function of the virtual target distance, is plotted in figure 4.13a; the five repeat trials to the same target distance are plotted as a single averaged data point.

During the experiment, it was observed that participant 4 was constantly throwing the real ball with much less force than the other participants. The participant had been observed successfully throwing 9 m distances with little effort required prior to the experiment, and appeared to be struggling to accept the safety of throwing the real ball so close to the screen. It was found that participant 4 accounted for the lower group of results in figure 4.13a, consistently throwing between 3-5 m distances for all target distances (with a standard deviation of less than 0.5 m for each group of trials). To prevent this lower group from distorting the results of the other three participants, only participants 1-3 will be considered when investigating the distance, as shown in figure 4.13b.

In figure 4.13c, the data from figure 4.13b are plotted as a series of linear best-fit lines labelled by configuration. This graph demonstrates much poorer scaling between the virtual distance and the thrown distance than we had expected to see, although the slopes do progress upwards slightly between 6 m to 9 m. Because the participants were throwing at targets suspended 1 m above the ground rather than a target placed on the ground plane, it is not necessary for the participant to throw the exact distance to hit the target, so throws travelling beyond the target distance for the closer targets may not present a problem. To look at the difference between the variables in greater detail, each of the three will be discussed in turn.

Figure 4.14a plots the distance thrown from position *A* (nearest the long edge of the CAVE) and figure 4.14b plots position *B* (nearest the short edge of the CAVE) against the target distance. The two variables appear to produce similar results, with a small number of throws from position *A* reaching farther distances than the average 6-9 m distance range. Using ANOVA on the two grouped sets of configurations 1-4 (position

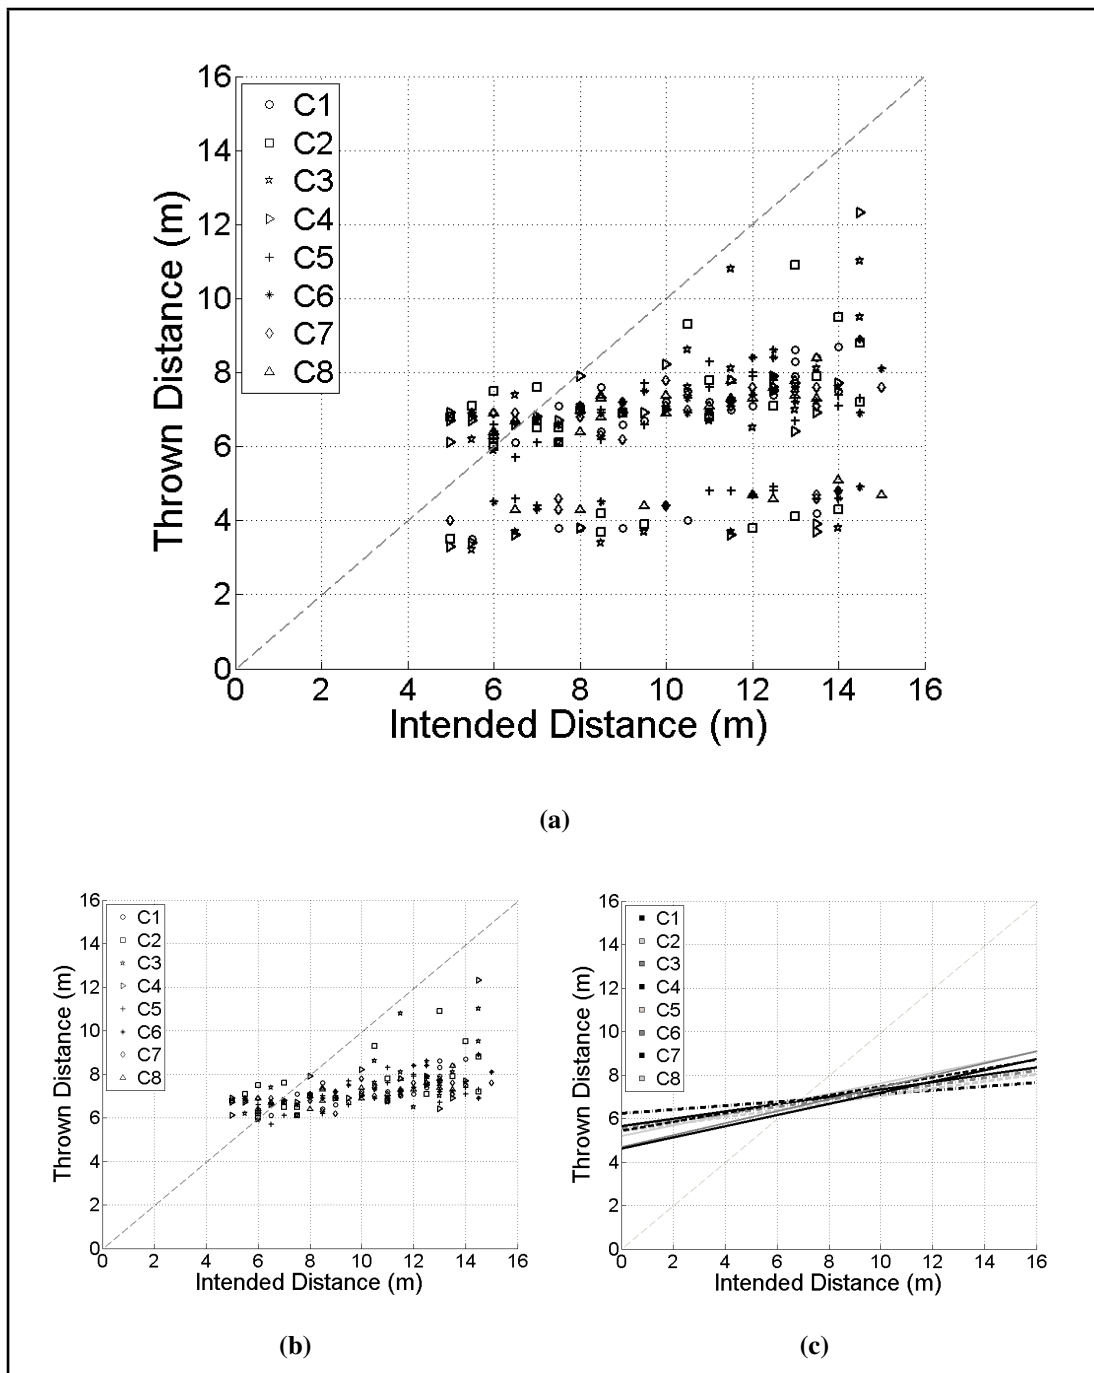


Figure 4.13: Each configuration is individually plotted (c1–8). The diagonal dashed line denotes the 1:1 ratio of the two axes. (a) For all participants, the average distance thrown as a function of target distance. (b) Participants 1–3, the average distance thrown as a function of target distance. (c) Participants 1–3, the average distance thrown as a function of target distance from figure 4.13b presented as a series of best-fit lines.

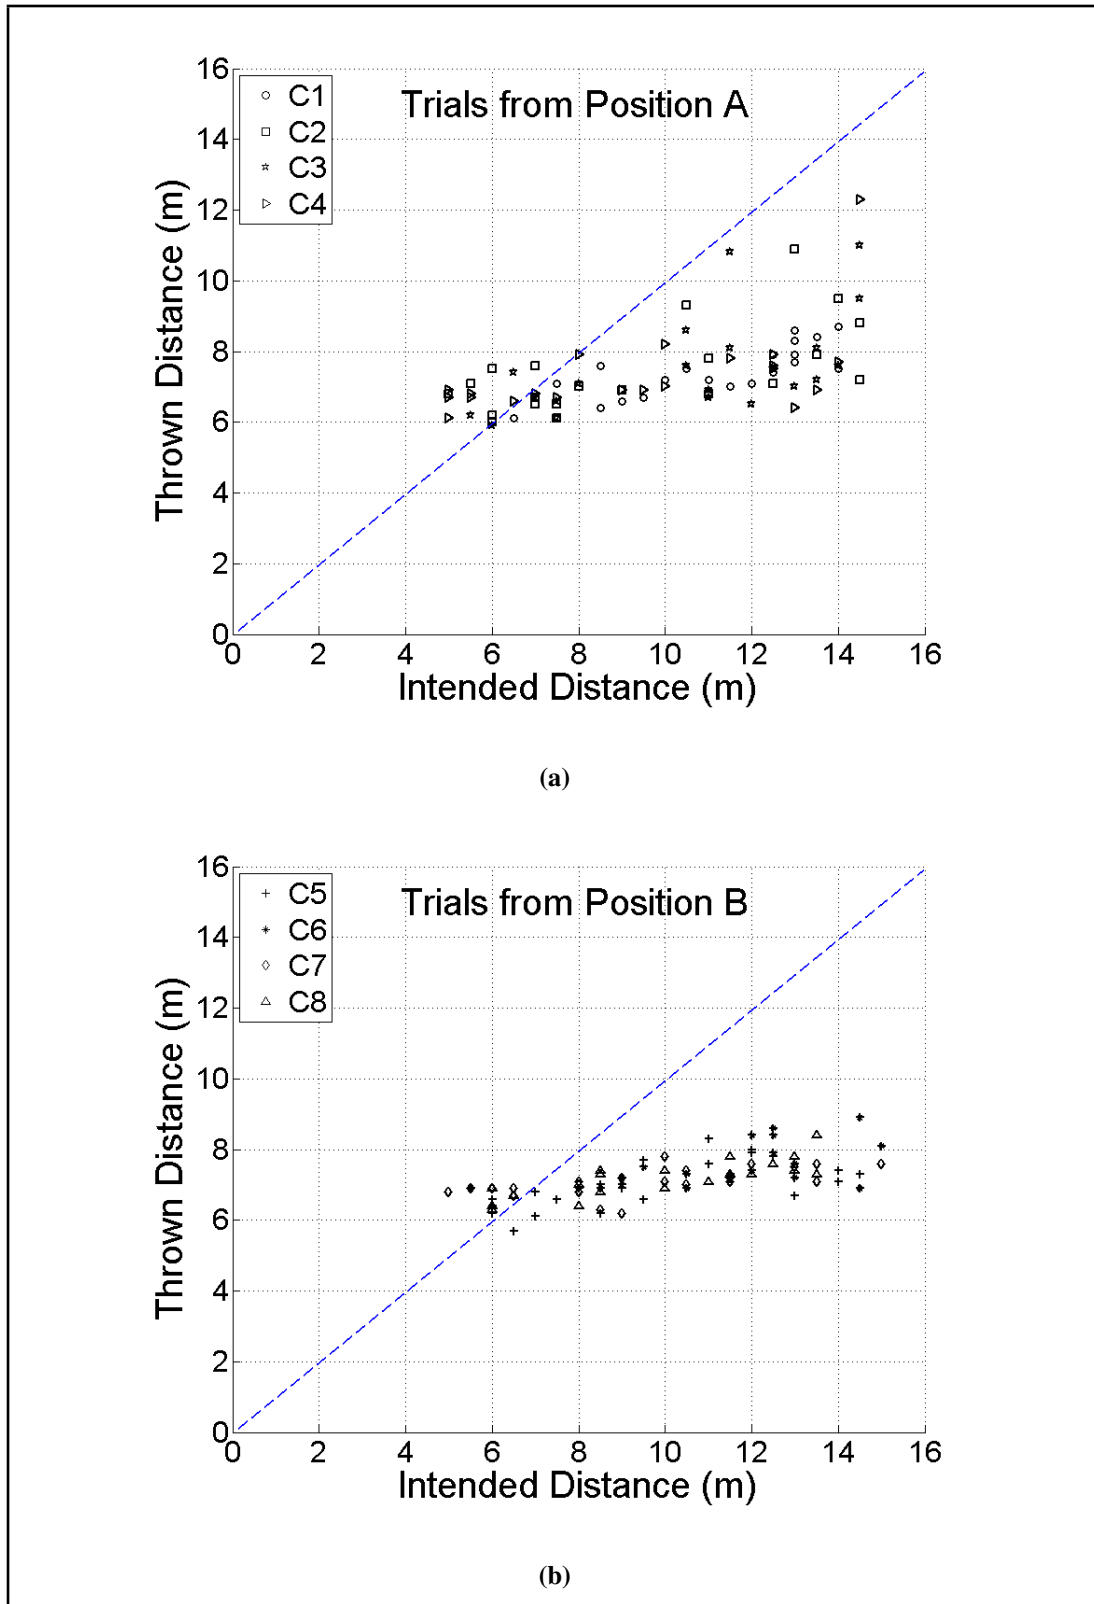


Figure 4.14: Average distance thrown by participants 1–3 across all configurations as a function of the target distance. Each graph only contains the data from the trials representing the variable condition designated in the title. (a) Configurations 1–4; (b) configurations 5–8. The dashed blue lines denotes the 1:1 ratio of the two axes.




Figure 4.15: Average distance thrown by participants 1–3 across all configurations as a function of the target distance. Each graph only contains the data from the trials representing the variable condition designated in the title. (a) Configurations 1, 2, 5 and 6; (b) configurations 3, 4, 7 and 8. The dashed blue lines denotes the 1:1 ratio of the two axes.

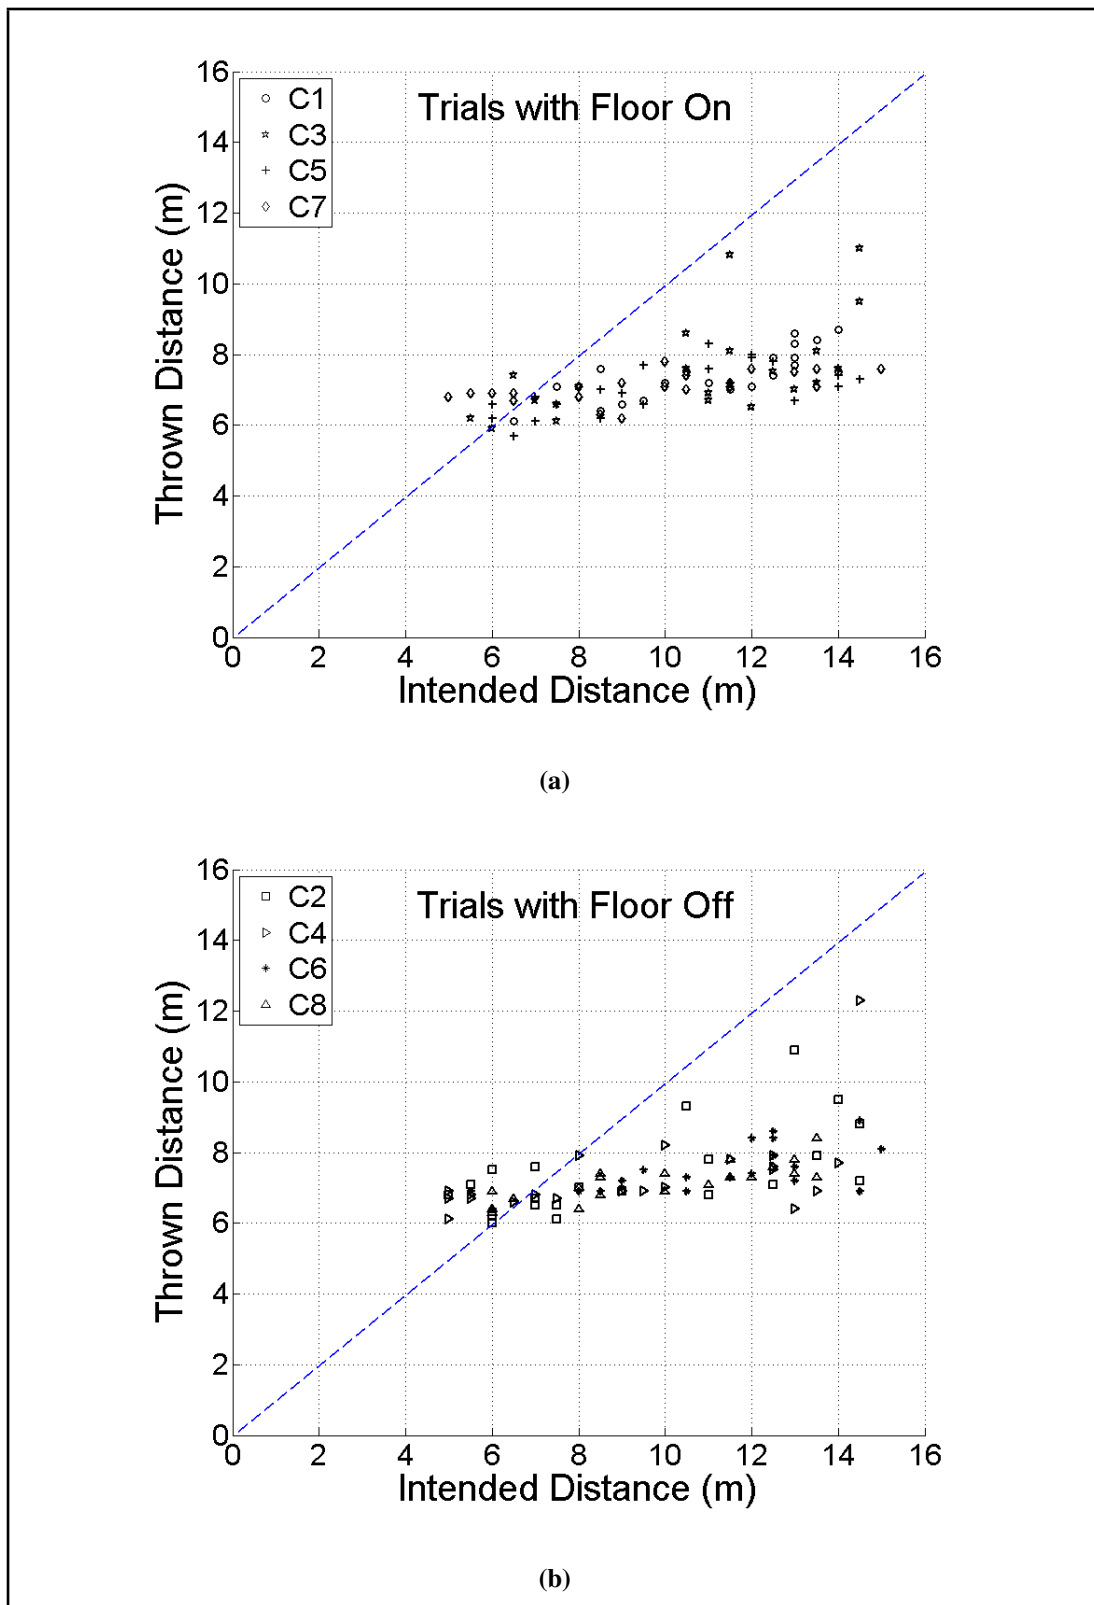


Figure 4.16: Average distance thrown by participants 1–3 across all configurations as a function of the target distance. Each graph only contains the data from the trials representing the variable condition designated in the title. (a) Configurations 1, 3, 5 and 7; (b) configurations 2, 4, 6 and 8. The dashed blue lines denotes the 1:1 ratio of the two axes.

A) and configurations 5-8 (position *B*) revealed that the null hypothesis of no difference in results between the two positions could not be rejected ($F(1, 68) = 0.22, p = 0.64$).

Comparison of stereo versus non-stereo is made in figure 4.15a and figure 4.15b, and floor versus no floor in figure 4.16a and figure 4.16b. Once again, ANOVA confirmed there to be no significant difference between the throw distances (stereo $F(1, 68) = 0.18, p = 0.67$ and floor $F(1, 72) = 0.14, p = 0.71$).

It was noted that there were a great number of trials where the ball missed the target to the left or right. Exploring this revealed that all participants threw too far to the left in 82% of the trials from position *B*, and threw too far to the right in 51% of the trials from position *A*. In this instance, participants 2-4 threw 66% of the trials to the right but participant 1 threw only 10% of the throws to the right. Exploring the statistics of the misses with ANOVA verified that there was a significant effect on the left misses from position *B* ($F(1, 40) = 108.3, p = 6.04E-13$), and right misses from position *A* ($F(1, 40) = 69, p = 3.07E-10$). Positioning in the CAVE was the only variable influencing this. The number of misses left and right were very evenly spread between the stereo on/off (left: $F(1, 40) = 0.34, p = 0.56$; right: $F(1, 40) = 0.17, p = 0.68$) and the floor projection on/off (left: $F(1, 40) = 0.87, p = 0.36$; right: $F(1, 40) = 1.01, p = 0.32$) variables.

Participant 1 had been in the CAVE previously for other studies. The difference in results, the lower number of misses to the left and right of the target compared to the other three participants is perhaps due to this previous experience. The other three participants had not undertaken any previous experiments in the CAVE used for the experiment, or any other to our knowledge.

The outcomes of the trials were analysed to explore whether they were affected by any design aspect of the experiment. Simplifying the categories to:

- ‘**ML**’ – Missed Left
- ‘**MR**’ – Missed Right

- ‘**OT**’ – On Target
- ‘**null**’ – for a failed capture

Participant 1 was found to be significant ($F(3, 1116) = 30.77, p = 0$), with **OT** throws mainly in trials by participants 1-3, **MR** and **null** by participants 2-4, and all participants had **ML** trials. The configuration was also found to be significant ($F(3, 1116) = 395.69, p = 0$), with **OT** trials mainly occurring in configurations 2-5. Configurations 2-4 were all from position A, but there are no common variables across configurations 2-5. Both the **ML** and **MR** were found mostly in target distances between 8-13 m, with **OT** between 6.5-12 m. Although this was found to be statistically significant ($F(3, 1090) = 8.8, p = 9.0991E-6$), it is not significant to the outcomes because the ranges cross over. Similarly, **OT**, **ML** and **MR** occurred most in the same range of target rotations, 0.1-0.4 radians; the difference in means was not found to be significant ($F(3, 1089) = 0.92, p = 0.4316$). Finally, the physical distance of the participant from the screen was found to have a significant effect ($F(3, 1089) = 28.51, p = 0$), where **OT** trials happened mostly when the participant was within a range of 1-1.6 m from the screen, **MR** when they were 0.5-1.5 m from the screen, **ML** at 0.7-1.7 m and **null** trials at 1-1.5 m. While the range of distances cross over, the **ML** and **MR** ranges fall closer than 1 m, suggesting that perhaps they are too close to the screen. If this were true, it would support the theory that the participants are fearful of throwing towards the screen, and may be aiming the ball slightly to the left or right to allow themselves more room to throw the ball.

Verbal feedback from the participants included two noting that the distances appeared correct to them, and one noted that he used the pitch markings in the virtual scene to aid him. Two of the participants stated that they noticed no difference when the variables were switched, while a third had noticed the floor being changed but not the stereo. Finally, two of the participants commented that it was “strange” and difficult to perform the task while being unable to see the virtual ball. While the version of

VERST not restricted for the experiment includes the virtual ball, it should be noted as there may be a difference when the virtual ball is visible. None of the participants noted any nausea or discomfort at any point during or after the experiment.

4.6 Discussion

In this preliminary study, it was found that the participants did not respond well to the perceived distance with most throws travelling between 6-9 m with little regard to the actual target distance. The use of stereoscopy and the use of a floor screen did not significantly affect the distance thrown by the participant. It was noted that the physical position of the participant and the position of the target they were aiming at may have influenced them to over-rotate their throw direction to the left and right. It is also possible that the participants were unintentionally aiming the throw closer to the diagonal across the floor screen – the longest physical space to throw. For this reason the L-shape configuration of the CAVE may not be appropriate to the task of throwing a rugby ball at a virtual target. In the future, these results should be compared with those with a flat and curved screen to confirm that the shape of the screen was the true cause.

One of the participants was unable to make any throws beyond 5 m, which nullified the training potential of the system in that case. It appeared that they were unable to accept the safety of the system, and were fearful of hitting the CAVE walls with the ball. It is conceivable that improving the sense of presence the system provides might assist the participant's performance. This demonstrates that the use of this VE for training will not be suitable for all users, as it is likely that some other users will experience similar problems in accepting the virtual content and ignoring the physical distance.

That two of the participants found it strange that they could not see the virtual ball

in flight suggests that they may have been attempting to merge reality and virtuality, but the lack of the virtual ball disturbed this.

The participants not always noticing a change in configuration settings may be attributed to inattentive blindness (Triesch et al. (2003); Cater et al. (2002), also known as task-blindness or perceptual blindness) in which their visual perception system was so concentrated on the task of throwing the ball that they failed to perceive change in the stereoscopy and floor screen variables. Inattentive blindness may also have contributed to the participants' feelings of strangeness in not being able to see the virtual ball. Similar to the concept of presence, the participants may have 'forgotten' they were not interacting with the real world and felt uncomfortable when the illusion was broken.

The system of fixing the ball to the ceiling of the room may have also caused problems for the system's functionality: not only was it unnatural to see the ball being pulled upwards by the rope, but all four participants concentrated on the movement of the ball after the throw, as it would swing backwards behind them and be difficult to catch again.

4.7 Summary

The performance of VERST in experiment 1 is briefly analysed in the context of the following research questions:

- What types of skills appear to be best suited to training in VEs?
- Which technologies achieve the best results in specific performance measures?
- Should stereoscopic displays be used?
- What does it cost?

The experiment revealed that the passing task may not be suitable for use as a training skill in VERST, as it is possible that the participants were not responding well

to that task of throwing the ball at a screen-projected virtual scene. The shape of the CAVE may have been causing participants to rotate their throws and aim towards the largest area of real free space. Stereoscopic presentation was not found to significantly affect the distance thrown, and may not be suitable for VERST. The equipment used in the experiment was very costly, and the experiment could not have been performed without the Visionair grant.

A number of potential issues and unexpected effects revealed themselves during this experiment, and the poor response of the participants to the virtual target stimuli may have been influenced by any or all of these factors:

- stereoscopy and the floor screen did not significantly affect the distance thrown;
- the direction of the participants' throws may have been influenced by the shape of the screen;
- the power of the throws may have been influenced by the task of throwing the ball towards the screen (fear of damage to the equipment);
- the participants were distracted by the swinging of the real ball after each trial.

To investigate this further, a second experiment was designed in an attempt to address some of these concerns by comparing different types of screens, using a different method of tethering the ball and further evaluation of the use of stereoscopy.

Chapter 5

Experiment 2: Further Evaluation of the Depth Perceived in VERST

A second exploratory experiment was designed to provide a further insight into the capabilities of the system whilst also attempting to investigate the system in the context of the following research questions:

- What types of skills appear to be best suited to training in VEs?
- Which technologies achieve the best results in specific performance measures?
- Should stereoscopic displays be used?
- What does it cost?

The main focus of experiment 2 was to investigate the use of different screen shapes, considering whether this might have affected the results of the throwing task. Participants performed the same task as in the first experiment, using both curved and flat screens. We were given access to a different facility than that used for experiment 1 (The Hartree Centre, 2014), with both high fidelity curved and flat display environments. However, because the tracking system available in the curved-screen laboratory only operates at 60 Hz, it would not provide sufficient tracking support for the throwing task. In order to use both the screens, a verbal estimation task was performed with

both configurations, and the throwing task with the flat screen only, with the intention of using the verbal results to draw comparisons between the curved and flat screens, and theorise about the performance of the participants in the throwing task using the curved screen. The results from the two experiments were not expected to be perfectly matched for each participant (Bridgeman et al., 2000; Parks, 2012; Goodale and Milner, 1992), but it is possible that some comparison between the results would have been found. Based on the review of distance perception experiments by Renner et al. (2014), the distances in the verbal estimation task were expected to be underestimated by approximately 26%.

This section details the porting of the system followed by the development and execution of the experiment.

5.1 Description of the Hartree Centre Facilities

The Hartree Centre contains two large power walls that are used by a variety of companies and academic researchers as virtual stages for many different purposes.

The curved screen (situated in the Leverhulme Laboratory, see figure 5.1b) was 2 m tall with an arc length of 11 m (cord length approximately 8 m). The screen was rear-projected by 8 separate projectors that ran on either a single unit or a cluster of 8 PCs, depending on the demands of the application. The screen had a resolution of 8500×1920 pixels. An ART tracking system consisting of 12 cameras was fixed around the edges of the screen. The cameras were able to track multiple objects in front of the screen at 60 Hz.

The flat screen (situated in the Virtual Engineering Centre (VEC) Laboratory, a facility run by Liverpool University; see figure 5.1a) was 6 m wide and 2.1 m tall with a resolution of 3387×1200 pixels from 4 projectors. The laboratory housing the flat screen was equipped with a 12-camera Vicon system that covered the entire floor area in front of the screen by placing the cameras high on the walls all around the room.

The cameras were able to capture movement at 250 Hz, making it more than capable of capturing high-speed ball movements.




Figure 5.1: (a) A virtual car cockpit, featuring a real car seat and steering wheel with a virtual exterior scene. (b) Testing the curved screen using the testbed environment.

The VERST system was implemented in the Unity and MiddleVR software packages and presented on the flat and curved power walls. Both screen configurations had stereoscopic projection, presented through tracked shutter glasses, which allowed the scene to be geometrically projected correctly for each user. Before the experiment, each participant's IPD was measured using a basic IPD ruler, and taken into account in the stereo projection.

The change of location for the second experiment meant that the real ball could not be suspended from the ceiling, so a bungee and wrist strap were used for the throwing experiment. The capture area for the ball was determined during each trial as being 1 m in front of the user (based on their position from the tracked glasses).

5.2 Experimental Design

Both the verbal and throwing estimation experiments would use two variables to create a testing configuration:

- whether stereoscopic presentation was active (S) or not (NS); and

- the physical distance of the participant from the screen (Near/Medium/Far).

The participant distances were originally designed to be at distances of 1 m (Near), 3 m (Medium) and 5 m (Far) from the centre of the screen, but were changed to 1 m, 3 m and 4 m due to space limitations in the laboratories.

For the verbal estimation experiment, this produced six configurations as shown in table 5.1.

Table 5.1: Trial configuration table for the verbal task in experiment 2.

Configuration Number	Stereo / Non-Stereo	Participant Distance from Screen
1	S	Near
2	S	Medium
3	S	Far
4	NS	Near
5	NS	Medium
6	NS	Far

Because the participants could not throw at close proximity to the screen, only the medium and far distances were used, reducing the throwing experiment to four configurations as shown in table 5.2.

Table 5.2: Trial configuration table for the throwing task in experiment 2.

Configuration Number	Stereo / Non-Stereo	Participant Distance from Screen
1	S	Medium
2	S	Far
3	S	Medium
4	NS	Far

To remove any bias in the data resulting from the order in which participants performed each experiment, the order in which the configurations were used was different for each participant (for full specification, see appendix 7.6.1), and the target distances were randomly generated in the range of 5-10 m.

Each participant would estimate the distance of every target (5, 6, 7, 8, 9 and 10 m) with three repetitions of each target rotated around the azimuth of the user (from the position of the tracked glasses). This meant each user would estimate 18 distances per configuration, twice for the verbal task and once for the throwing task, resulting in:

$$18 \text{ distances} \times (6 \times 2 \text{ verbal} + 4 \text{ throwing configurations}) = 288 \text{ trials per participant.}$$

5.3 Developing the Experiment

To port the system for use at the Hartree Centre, VERST was created using the Unity software package. Unity allowed the fast development of the verbal and throwing tasks for the experiment in high quality to enhance the fidelity.

MiddleVR was used to facilitate connections between the hardware and software, as it allowed flexible functionality with VRPN, and provided a full visual interface to aid the process of setting up a new VE.

For tracking the ball with the accuracy provided in CRVM, both spherical and flat, circular markers were used. The spherical markers were attached to the ball with Velcro® pad, so that they could be quickly replaced without difficulty. Flat markers were used to the rear of the ball so that they would not interfere with the subject holding the ball.

5.3.1 Porting of the System to Unity

First, the static scene was created in Unity, with interaction being added later. The Millennium Stadium model (figure 5.2a) used in experiment 1 (Model created by user ‘dizzyHARSH’, 2014) was downloaded in COLLADA format and imported into Unity.

The model had been created in the style of a football pitch, requiring editing of the pitch markings and the goals to transform it into a rugby pitch. Import errors in the model ground plane (figure 5.2b) forced the creation of a new pitch floor: two

planes were added and placed a few units above the original ground plane to mask it; one plane was textured with a repeating grass image (figure 5.2c), while the other was textured with pitch markings where the texture was transparent around the lines. The fully-textured pitch is shown in figure 5.2d.

The pitch model had been created as a single complex mesh, and so the football goals were removed by texturing them with a completely transparent material. The rugby posts were created from 5 primitive cylinders in 3D Studio Max – three formed the bars and two formed the outer pads at the base of the vertical posts (see figure 5.3a). While the posts were created using the correct IRB dimensions in 3D Studio Max, they were positioned so far from the user in the VE that they suffered from flickering due to *z*-buffer conflicts (*z*-fighting). To counter this (which was found to be painful for the eyes during testing at the Hartree Centre), the posts were all thickened (depth and width were extended) to be much larger than is natural (see figure 5.3b). From the position of the participant, this resolved the *z*-fighting without the posts seeming disproportionately large.

Next, a bull's-eye target was created in 3D Studio Max from a flattened cylinder and three boxes (see figure 5.3c). In Unity, bump maps were added atop the wood and target (Barbati, 2014) textures to enhance the look of the target. A Blob Shadow Projector was placed in the centre of the target to cast a shadow to the ground below, also greatly improving the fidelity of the scene.

Finally, a ball object was created from a deformed sphere in 3D Studio Max (figure 5.3d). The ball was set up as a prefabricated item ('prefab') in Unity, to allow it to be called as a fresh object for each trial. The completed static scene from the viewpoint of the user can be seen in figure 5.4.




Figure 5.2: (a) and (b) The original stadium model. (c) The grass image used to texture the ground of the pitch. (d) The full pitch environment.

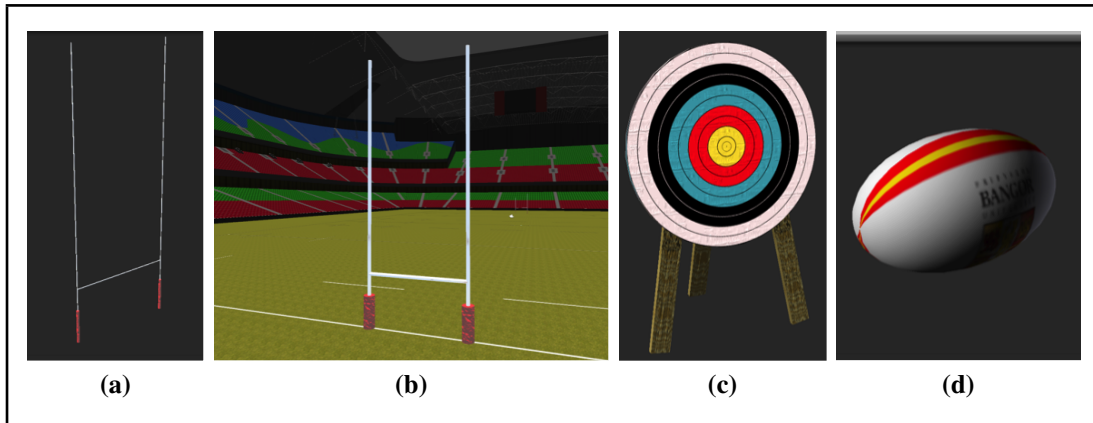


Figure 5.3: (a) The original goal post model; the posts are sized based on the IRB regulations. (b) Re-sized posts to counteract the z-fighting. (c) The textured target model in Unity. (d) The textured ball model in Unity.




Figure 5.4: The scene from the viewpoint of the user.

5.3.2 Scripting Interaction

Both the verbal and throwing experiments had very similar interaction, scripted using C#. Looping through each configuration, the list of target distances (three repeats of each whole metre from 5–10 m) is shuffled. Each trial is loaded consecutively, the target is positioned and rotated, and a timer started. For the verbal task the timer waits the specified time then hides the target and begins the next trial. No further action is required as the verbal estimate is recorded by an observer. For the throwing task, the timer is started and the ball tracked. Using the same system of data capture detailed in chapter 3, once the ball passes a specified position, the data is processed and the flight path calculated.

The ordering of the configurations was specified in a series of accompanying text files which could be used to make some experimental design changes to the program after it had been compiled (see appendix 7.6.2 and appendix 7.6.3).

When a new configuration was loaded, the two variables of stereo on/off and the participant positions were modified for the given configuration. For stereo-enabled configurations, the distance between the stereo cameras within Unity were set to be the subject's IPD so that each eye would receive an image correctly parted from the other eye. When stereo was disabled, the distance was set to be zero so that the each eye would see the same image. The participant would receive on-screen instructions about which position to stand at (designated 'Nearest' 'Middle' and 'Furthest'), and would be given the specified number of seconds to move (see figure 5.5).

The target distances were held as a list, shuffled at the beginning of each configuration using the Fisher-Yates inside-out shuffling algorithm. The algorithm was used for its ability to produce unbiased shuffling, with every permutation equally probable.

The experiments produced data files containing the target distance (the magnitude and the absolute position in the environment), the rotation of the target around the user and the head position of the user for each trial. Figure 5.6 is an example of one of the files from the verbal experiment, while the throwing experiments also recorded the




Figure 5.5: The instruction text and countdown displayed in the centre of the screen.

Participant 1, Configuration 1				
Configuration started at 9/13/2013 10:12:57 AM				
Trial No.	Target Pos.	Target Dist.(m)	Target Rot.(deg)	Head Pos.
1	(7.7, 0.0, 2.2)	8	-16	(0.1, 1.8, -0.2)
2	(9.9, 0.0, -2.1)	10	12	(0.0, 1.8, -0.4)
...				
18	(7.8, 0.0, -2.4)	8	17	(0.0, 1.8, -0.5)
Configuration completed at 9/13/2013 10:14:40 AM				

Figure 5.6: Example output file of the verbal task in experiment 2.

position of the ball and the average velocity recorded for the physics engine, followed by the ball's point of impact on the ground of the pitch.

The average velocity of the ball was obtained using an average of six points converted into five velocities. The data-collection function would begin collecting data as soon as the trial started, and once six points had been collected, it would replace the old data first – always maintaining the data of the last six points it passed through. This data would be collected until the tracked ball moved a certain distance (specified

Table 5.3: The MiddleVR settings for the curved and flat screens in experiment 2.

	Curved Screen	Flat Screen
Ball Tracker Address	-	ball@VEC-VICON
Head Tracker Address	DTrack@192.168.0.19	head-theia@VEC-VICON
Screen Node:		
Position	[0.000,-3.000,1.194]	[0.000,-1.480,1.140]
Width (m)	10.249	6
Height (m)	2.387	2.1
Viewport Node		
Width (pixels)	8025	3387
Height (pixels)	1920	1200
Stereo Camera Node		
Vertical FOV	60	60
Near Clipping Distance	0.1	0.1
Far Clipping Distance	1000	1000

in the settings text file) away from the head position of the user. This would allow the collection of the last six points before the ball was released by the user, providing the ideal data to the physics engine.

5.3.3 Hardware Interaction with MiddleVR

MiddleVR provide a way of easily connecting devices and a sophisticated way to use the hardware in VEs. It also allows the configuration of the hardware after the VE has been compiled, providing the necessary links to the software components are set up pre-compilation.

The node hierarchy in the 3D Nodes view was built up to cater for the hardware and VE. For VERST, one tracker was attached to the stereo camera so that the viewpoint moved with the glasses, while another tracker was attached to the ball. The different settings for the two experiment measures and the two screen types are summarised in table 5.3.

5.4 The Experiment

5.4.1 Experimental Procedure

Before the participant arrived for the experiment, the system was prepared. Both the flat and curved screen room systems were initialised and the projectors allowed time to warm up. The tracking systems were not re-calibrated prior to each experiment, as the participants were all scheduled for the same day and the systems are re-calibrated regularly as part of the centre's maintenance plan. The system was briefly checked before the experiments began, to ensure regular operation.

When each participant arrived, they were greeted and provided with an information sheet describing the experiment (see appendix 7.6.4). Once they had read the sheet, they were given the opportunity to ask any questions, and were asked to complete the consent form.

The participants undertook a short stereoacuity test, to ensure their stereoacuity falls within the normal range of 20–500 arcseconds (0.3–8.3 arcminutes). The equipment used for the experiment at CRVM belonged to Aix-Marseille University. For this experiment, a simple test provided by the Department of Psychology at Bangor University was used. The participant was required to view three random dot stereograms using anaglyphic glasses (see figure 5.7). In figure 5.7a and figure 5.7b, the participant was informed they should be able to see a square and a spiral emerging from the image; then they would be asked to describe what they could see in figure 5.7c. Figure 5.7c contains a sinusoidal corrugation pattern, with depth in the peaks and troughs; the peaks and troughs are oriented at approximately 30° from the horizontal. Participants without stereoacuity in the normal range would not be able to see this pattern.

Once the stereoacuity test was complete, the participant's IPD was measured using an IPD ruler (see figure 5.8). The ruler is placed with the curved cut-out section resting on the subject's nose, with the rectangular cut-out sections in front of each eye. The IPD is measured by determining the measurement at the centre of the pupil for each

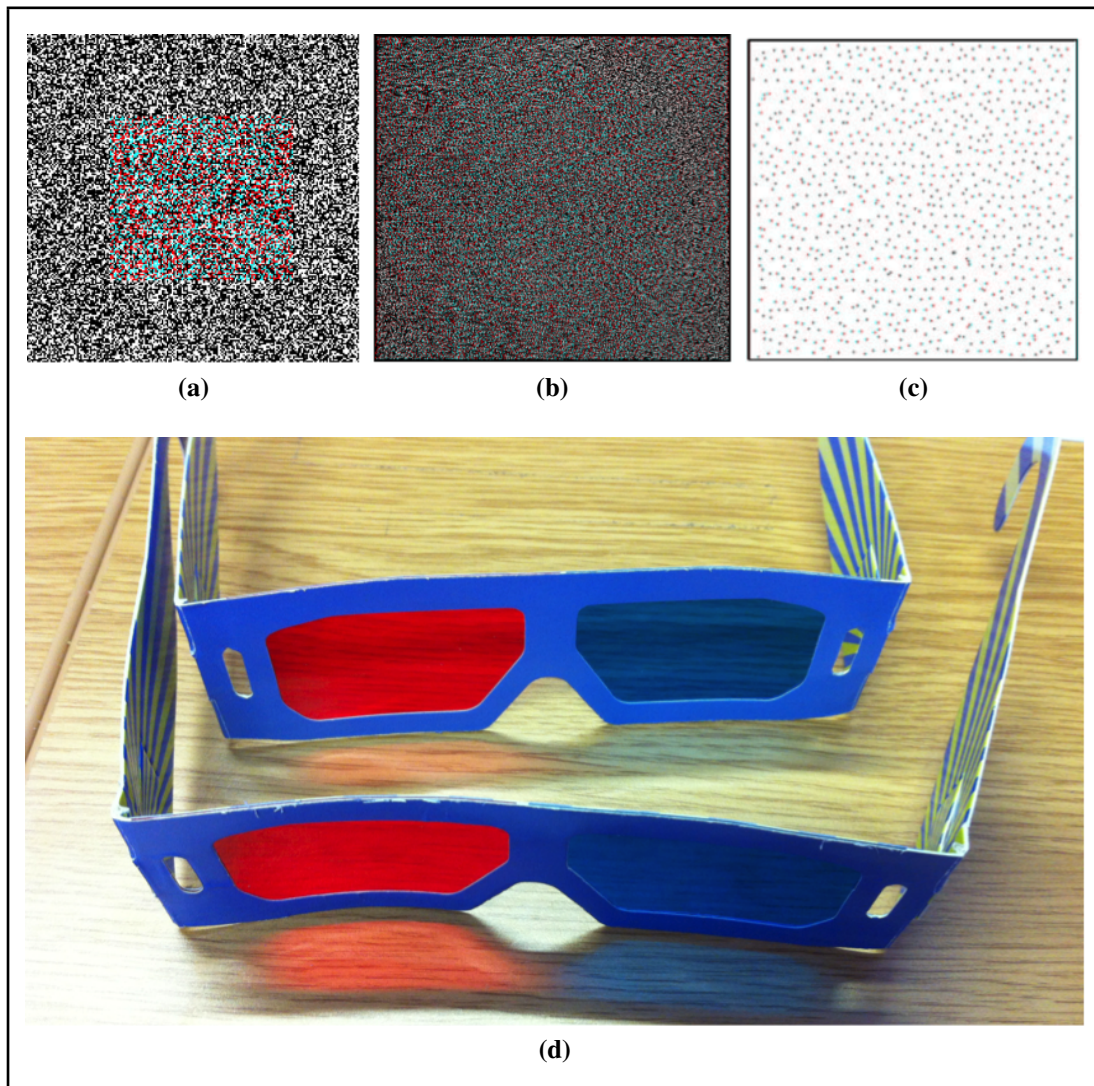


Figure 5.7: (a), (b), and (c) are random dot stereograms to be viewed with anaglyphic glasses (d).




Figure 5.8: The IPD ruler: to use, hold to subject's face and measure the centre for each pupil eye (the rectangular ruled sections represent the distance from the centre of the ruler), then adding the two separate measurements provides the distance between the pupils. The IPD was used to configure MiddleVR's stereoscopy, to ensure the separation of the two images was correct for each participant.

The participant was then allowed time to practise throwing with a ball of the same construction as the one used in the experiment – a Skill-tec size 4 training ball on a bungee cord attached to a wrist strap. The participant was observed to ensure their throwing skills were sufficient, and they had the opportunity to familiarise themselves with the wrist strap and bungee.

The rooms in which the task would be performed first were alternated for each participant. When the participant was ready, they were lead to the first room and performed the verbal task. If they were in the Leverhulme laboratory, they then moved straight to the VEC laboratory, but if they were in the VEC laboratory, they remained to perform the throwing task before moving to the Leverhulme laboratory.

For both tasks, the participant would stand at the instructed position for each configuration of both tasks. A bull's-eye style target would appear in the virtual playing field in front of them. They would have 5 seconds to react to the target by either stating the target distance or throwing the ball at the target. For the verbal task, the participants were asked to provide an answer in whole metres. The participants were not given any information about the range of target distances, and any comments or

questions regarding the distance of the targets or the participant's three positions were neither confirmed nor denied by the experimenter.

During the experiment, feedback was noted from the participant. After the experiment was completed, the participants completed a Presence Questionnaire to evaluate their immersion (see appendix 7.6.5).

In figure 5.9a, a participant is performing the throwing task, while in figure 5.9b another participant is undertaking the verbal task.




Figure 5.9: Participants during the experiment performing the (a) throwing and (b) verbal tasks.

5.4.2 Participants

Six participants aged 20–54 years took part in experiment 2. For this experiment, the participants were not required to have extensive previous experience playing rugby; five of the participants had played only at school or college level, while one had played for a local team. Five of the participants were right-handed, while the sixth was left-handed; all reported normal or corrected-to-normal vision. Participants were all found to have stereoacuity in the normal range using random dot stereogram. The details of the six participants are tabulated below in table 5.4, followed by details and observations recorded about each participant during the experiment.

Table 5.4: Details of the participants recruited for experiment 2.

Participant ID	Age	Gender	Years playing rugby	Experience (position of play)	Left/right-handed	Wore glasses/contact lenses	Other eye problems	Colour-blind	Stereoacuity in normal range	Consented to being filmed	Date of participation
1	30	M	-	School Only	R	Glasses	N	N	Y	Y	13/09/13
2	54	M	-	School Only	R	Glasses	N	N	Y	Y	13/09/13
3	48	M	-	School Only	L	N	N	N	Y	Y	13/09/13
4	54	M	-	School Only, centre	R	N	N	N	Y	Y	13/09/13
5		M	-	Second row, secondary and college teams; both union and league	R	N	N	N	Y	Y	13/09/13
6	20	M	5	Flanker for Helsby RFC	R	N	N	N	Y	Y	13/09/13

Though the experiment had been designed to test participants who regularly played rugby and had extensive experience of the sport (as with the subjects who participated in experiment 1), the group of recruited participants failed to attend on the morning of

the experiment, which led to the use of participants who had only a small amount of previous rugby experience (except for participant six). The majority of the participants had only experience of playing rugby while at school, with many years passing since they last played.

An additional problem encountered during the experiment was that two of the participants failed to complete all three sections of the experiment due to time constraints:

- participant 1 completed the throwing and the presence questionnaire, but only performed the verbal task in the flat-screened laboratory;
- participant 2 completed the verbal and throwing tasks, but not the presence questionnaire.

Fortunately, the two sections that were not completed were different; this resulted in five complete sets of results for the verbal experiment, five completed presence questionnaires and six complete sets of results for the throwing experiment.

Participant 1 often waited until the target had disappeared before providing a response in the verbal exercise, and appeared to be comparing pairs of targets.

Participant 2 referred to the participant position marker line at 4 m as being 5 m (he was not corrected), suggesting that he underestimates real distances. During the throwing task, he threw the ball softly but took a large step forward each time. His verbal estimations were very confident.

Conversely, participant 3 was very unsure when providing distance estimates. He stated that he found distance estimation easier when using the curved screen.

Participant 4 was aware of the pitch markings, and was granted his request to use them to aid his judgement. However he believed the 10 m line to be 20 m away from himself, making the cue a hindrance rather than a help (again, he was not corrected). He suggested it would be better to have a virtual rugby player to throw the ball to. Although this was not implemented in the system for this experiment, it is considered as future work to be done on the system.

Participant 5 very quickly stated that he found the verbal task difficult, saying that he used points of reference rather than intuition to tell the location of objects. Despite this, in conversation he correctly identified the furthest position as being 4 m from the screen without any known references. He was clearly uncomfortable with the VE, and did not enjoy any aspect of the experience. He also stated that he found the curved screen to be no better for distance perception than the flat screen, and that the stereoscopic glasses interfered with the throwing task. He often appeared frustrated with the task and used self-deprecating language, saying repeatedly that he ‘couldn’t get it right’.

Participant 6 was the most experienced, and the youngest participant. He had a much stronger throw than the other participants. He stated that he found himself concentrating on the target, and found the task absorbing. He also noted that after a few trials, he learned that shifting his body position slightly from side to side helped him to determine the distance of the target (he used motion parallax as a depth cue).

5.4.3 Data Collected

For all of the trials in both of the verbal and throwing experiments, the following data were collected:

- position of the participant from the tracked glasses;
- position of the target in the virtual space and the distance from the participant;
- the amount (in degrees) the target was rotated around the participant;
- the start and completion timestamps for each configuration and the whole experiment.

For each of the trials in the verbal experiment, the distance estimate by the participant was noted as the experiment progressed. As the target distances were shuffled

prior to each participant performing the experiment, the experimenter only knew the range of possible distances.

For each of the trials in the throwing experiment, the following data were collected:

- the position and velocity vectors that would have been input for the physics engine;
- the position the ball would land at according to the physics engine.

After both the verbal estimation experiment and the throwing experiment, the participants each completed a copy of the Witmer & Singer Presence Questionnaire (Witmer and Singer, 1998). The questionnaire was chosen as it has been used extensively in the evaluation of virtual environments since its creation in 1998. It has since undergone revision in 2005 (Witmer et al., 2005). The questionnaire is based on the semantic differential principle. Comprised of 24 seven-point items, it rates the participants' attitudes towards seven aspects of the VE using a Likert Scale format:

1. realism;
2. the possibility to act in the virtual world;
3. the quality of the interface;
4. the possibility to examine the virtual world;
5. a participant's self-evaluation of their performance inside the VE;
6. sounds used in the VE;
7. haptics used in the VE.

The ball itself is a haptic device, being a physical object with which the user is able to interact with the VE. The VERST VE did not, in this implementation, use any sounds, and so group six (items 20, 21 and 22) are discounted from the questionnaire results.

The questionnaire is analysed by adding the item results, as follows:

1. realism = Items 3 + 4 + 5 + 6 + 7 + 10 + 13.
2. possibility to act = Items 1 + 2 + 8 + 9.
3. quality of Interface = Items 14 + 17 + 18, all with scores reversed.
4. possibility to Examine = Items 11 + 12 + 19.
5. self-evaluation of Performance = 15 + 16.
6. haptics = Items 23 + 24.

The results of the presence questionnaire should only be used in the context of support for the verbal and throwing experiment data. Insight into the attitudes of the participants can be useful in understanding their results (Slater and Garau, 2007).

5.5 Results

5.5.1 Verbal Task

First, the individual target distances were analysed (across all participants and configurations, with both the flat and curved screens). The results are plotted in figure 5.10 and shown in table 5.5. Medians and inter-quartile ranges are used instead of mean and standard deviation because the distance estimates were made in whole metres.

Figure 5.10 displays the results of the verbal task in the Tukey boxplot format, where each box represents the data obtained for that intended target distance. The lower edge of each box represents the first quartile (above which lies 75% of the results), while the higher edge represents the third quartile (above which lies 25% of the results); the line dividing the box is the second quartile (the median of the data). The




Figure 5.10: Results of the verbal task. Intended perceived distance between 5–10 m is plotted against the distance estimated by all participants. The dashed line represents the 1:1 ratio of the two axes.

whiskers of each box represent the remaining data that falls within 1.5 times the inter-quartile range (the third quartile minus the first quartile), with the remaining points individually plotted.

The data in figure 5.10 shows that the target distances were mostly overestimated, target distances of 5, 6, 8 and 9 m only having the lower 25% of the results underestimated. For 7 and 10 m target distances, the lower quartile fell 1 m under the distance they were intended to perceive. The median results for 6 and 7 m targets are 3 m

Table 5.5: Descriptive statistics of the verbal experiment, plotted in figure 5.10.

Intended Target Distance (m)	5	6	7	8	9	10
Median	7	9	10	10	11	12
Maximum Estimate	24	22	22	24	25	27
Minimum Estimate	1	1	1	1	2	2
Q1	5	6	6	8	9	9
Q2	10	11	12	12	14	15
Inter-Quartile Range	5	5	6	4	5	6




Figure 5.11: The results of the verbal task in experiment 2 with the modified results of participant 4. The dashed line representing the 1:1 ratio of the intended to estimated distances.

above the intended distance, while the 5 and 8–10 m targets were 2 m above. The inter-quartile range indicates that 50% of the estimates were within a range of up to 6 m across all target distances.

Because of participant 4’s mistake in believing the 10 m pitch line to be the “20 m line”, it was considered possible that he had made his estimates based on a misconception, and that all of his data might be 10 m further than he had perceived. While this cannot be proven, the target distances were re-examined with his modified data (see figure 5.11). It is possible to see that any potential mistakes in his results have not greatly skewed the overall outcome of the data.

2×2 repeated measures ANOVA was conducted on the mean estimated distance values of the verbal task, revealing that the use of stereoscopy did not significantly affect the estimates made using either the curved ($F(1, 168) = 2.14, p = 0.15$) or flat ($F(1, 204) = 0.24, p = 0.63$) screens; the combined mean estimated distances were




Figure 5.12: Results of the verbal task separated into the estimations made under different configurations (refer to table 5.1 and table 5.2 for configuration descriptions) using the curved and flat screens. The horizontal dashed lines mark the range of target distances intended to be perceived (5–10 m).

also not significantly affected by the use of stereoscopy ($F(1, 348) = 0.48, p = 0.49$) (see figure 5.13a and figure 5.13b).

The position of the participant relative to the screen (at Near, Medium and Far distances using both the curved and flat screen) was found to have a significant effect ($F(2, 342) = 28.18, p = 0$) on the estimated distances; target distances were perceived as nearer the closer the participant was to the screen (see figure 5.14a, figure 5.14b and figure 5.15). The position of the participant was found to significantly affect the estimated distance when using the curved screen ($F(2, 162) = 80.25, p = 0$), but not the flat screen ($F(2, 198) = 1.9, p = 0.15$) (see figure 5.12).

Whether the estimates were made using the curved or flat screen was found to have a significant effect ($F(1, 348) = 40.4, p = 0$), with mean estimates made using the flat screen consistently overestimated 2–3 m farther than those made using the curved screen.

Although the target distances were measured from the head position of the user




Figure 5.13: Average distance verbally estimated by participants 1–6 across all configurations as a function of the intended target distance. Each graph only contains the data from the trials representing the variable condition designated in the title: (a) configurations 1–3; (b) configurations 5–6. The dashed pink lines denotes the 1:1 ratio of the two axes.




Figure 5.14: Average distance verbally estimated by participants 1–6 across all configurations as a function of the intended target distance. Each graph only contains the data from the trials representing the variable condition designated in the title: (a) Configurations 1 and 4; (b) configurations 2 and 5. The dashed pink lines denotes the 1:1 ratio of the two axes.




Figure 5.15: Average distance verbally estimated by participants 1–6 across all configurations as a function of the intended target distance. Each graph only contains the data from the trials representing the variable condition designated in the title: Configurations 3 and 6. The dashed pink lines denotes the 1:1 ratio of the two axes.

and maintained in the 5–10 m range throughout, it appears that the real space affected the participants perception of distance. For the targets, evenly spread in the range of 5–10 m, the ideal mean distance would be 7 m; removing the physical distance from the Near, Medium and Far positions places the mean estimation at 7 m (see table 5.6).

Table 5.6: Median verbal estimations by participant position, minus the physical distance between the participant and the screen when they are in that position.

Participant position	Median estimation (m)	Median estimation minus physical distance from screen (m)
Near	8	$8 - 1 = 7$
Medium	10	$10 - 3 = 7$
Far	11	$11 - 4 = 7$

Re-analysing the data minus the physical distance of the user at the time of the




Figure 5.16: The estimated distance minus the physical distance of the participant from the screen is plotted against the intended perceived distance. The dashed line representing the 1:1 ratio of the intended to estimated distances.

estimation changes the descriptive statistics (see table 5.7). The median values each decrease to 1 m under their respective intended target distances, and the inter-quartile range for each intended target distance becomes 4. The data is plotted in figure 5.16, the dashed line representing the 1:1 ratio of the intended to estimated distances.

Table 5.7: Descriptive statistics of the verbal experiment: comparing the original data to the data minus the participant's physical distance from the screen.

Intended Target Distance (m)	Data Minus Physical Distance						Original Data					
	5	6	7	8	9	10	5	6	7	8	9	10
Median	4	5	6	7	8	9	7	9	10	10	11	12
Maximum Estimate	12	13	15	16	18	18	24	22	22	24	25	27
Minimum Estimate	1	1	1	1	1	1	1	1	1	1	2	2
Q1	2	3	4	5	6	7	5	6	6	8	9	9
Q2	6	7	8	9	10	11	10	11	12	12	14	15
Inter-Quartile Range	4	4	4	4	4	4	5	5	6	4	5	6

5.5.2 Throwing Task

The mean results of the throwing experiment are plotted in figure 5.17. The results do not appear to be scaled to the intended target distances, with results ranging from 2–30 m regardless of the target distance.

2×2 repeated measures ANOVA was conducted on the mean thrown distance values of the throwing task, showing no significant effect from the use of stereoscopy/no stereoscopy ($F(1, 132) = 0.4, p = 0.53$) or the distance of the participant from the screen (Medium and Far distances away, $F(1, 132) = 2.18, p = 0.14$)

5.5.3 Presence Questionnaire

The results of the Presence Questionnaire are summarised in table 5.8, with the full results in appendix 7.6.6. No group had significantly different results to the other four, $F(4, 25) = 0.91, p = 0.47$, although the Possibility to Examine group has a lower score overall than the other four groups.

Table 5.8: Presence Questionnaire Group Results

Group	Maximum Possible Score	VERST Median Score	VERST Inter-Quartile Range
Realism	49	37	3
Possibility to Act	28	23	4
Quality of Interface	21	11	4
Possibility to Examine	21	17	2
Self-Evaluation of Performance	14	13	4
Haptics	14	10	1
Total	133	101	7




Figure 5.17: Distance thrown according to the results of the physics engine plotted against the intended target distance. The dashed line represents the 1:1 ratio of the two axes.

5.6 Discussion

5.6.1 Comparing the Results of the Verbal and Throwing Experiments

In terms of execution, the second experiment was performed in a similar manner to experiment 1. Differences included the use of the wrist strap rather than a ceiling tether, and a combination of spherical markers attached using hook and loop tape and flat, circular markers for tracking the ball. The wrist strap proved to be difficult to use, like the ceiling tether, resulting in the participants concentrating less on the task and more on the potential of the ball to hit them during the bungee's post-throw reflex. While the spherical markers were held solidly during most of the trials, as the ball hit the ground, the markers would occasionally be knocked off.

The participants responded to the verbal experiment with better scaling to the intended target distances than the throwing experiment. The verbal estimates were generally overestimated by participants, which conflicts with previous similar experiments (Renner et al., 2014), likely because of the different scenes used for the experiments. The curved and flat screens were found to affect the verbal estimates, with the flat screen estimates tending to be higher than the curved screen. The use of stereoscopy was not found to have a significant effect on the estimations being made in either the verbal or throwing experiments. A significant effect of participant's physical distance from the screen was found for the verbal estimation task, but not the throwing task. The results of the Presence Questionnaire showed that the participants had a positive attitude towards their immersion experience in VERST.

Comparing the results of the verbal and throwing tasks in the second experiment (see figure 5.18) showed an effect of the measurement type, where verbal estimates were tending to be larger than throws. Verbal estimates were scaled to the intended target distance, although they were overestimated, while distances thrown displayed no scaling.




Figure 5.18: Estimated distance during the verbal (dashed line) and the throwing (solid line) experiments plotted against the intended distance. The lighter-coloured dashed line represents the 1:1 ratio of the two axes.

5.6.2 Follow-up Questions for the Participants

As the results of experiment 2 had failed to clarify the results of the first experiment, the participants from experiment 2 were contacted to request further information about their experience of using VERST. The participants were asked the following questions:

1. Do you have any previous experience with a virtual environment?
If so, please give details.
2. How often do you play computer games?
[never - not often - often - very often - a lot]
3. Were you afraid of causing any damage to the screen or equipment by throwing the ball?
4. Did you find the task of throwing the ball at a target difficult or unnatural?

The participant's responses can be seen in table 5.9.

Table 5.9: Participant's responses to the four follow-up questions asked to them.

ID	Question 1	Question 2	Question 3	Question 4
1	Yes ^a	Not often	I was confident no damage would be done	No it was good
2	I dont have any precious experience of using virtual environments. I have seen simulators on television for aircraft, trains and formula 1 etc.	Not often	Initially yes but as time went on I realised that I had good control of the ball. Or at least I thought I did.	I didnt find the task difficult. However, I did find it slightly unnatural as in the real sporting environment the passes particularly the longer passes wold be directed at a moving target. Or where the target would be in a split seconds time.
3	NO	Not often	NO	NO
4	NO	Not Often	NO	A LITTLE UNNATURAL
5	No previous experience in such a virtual environment	Very often	Nope as the ball was tethered to my wrist	I found it difficult as I was stationary and not running forward as I would be if I was actually playing rugby
6	No, nothing worth mentioning other than the odd 3d movie	Often	At first yes because I didn't have a sense of the elastics stretching ability, but that caution went after I got used to it	It wasn't difficult, but it was unnatural because I am used to putting as much spin on the ball as I possibly can, but with this, when I originally tried to spin it how I would with a real ball, the computer couldn't work it out, so I had to do more of an under-armed toss with very little spin.

^aFrom working at the VEC Laboratory

Questions 1 ('Do you have any previous experience with a virtual environment?') and 2 ('How often do you play computer games?') were intended to establish whether the participants may have been overwhelmed by the equipment used in the experiment. The responses show that there was a mix of those who have had no experience in VEs, those who were aware of the technology and those who had a great deal of experience. Similarly there was a mix of participants who play games very often, to not very often at all (though notably, none said they 'never' play games).

The third question – 'Were you afraid of causing any damage to the screen or equipment by throwing the ball?' – was asked to verify the hypothesis to explain the results from experiment 1. The responses indicate that the initial hypothesis was incorrect. Participant 6 indicated that he was cautious at first, but was then comfortable; the other participants reported feeling no fear of damage.

The fourth question – 'Did you find the task of throwing the ball at a target difficult or unnatural?' – was asked to investigate whether the poor results may have been a result of the task type. Only one participant noted difficulty with the task due to being stationary as opposed to running (as would be done normally in rugby passing), while the others reported no difficulty. Three of the participants stated they found the task unnatural, with two of those indicating it was because they were unable to act as they would in the equivalent reality.

Because the responses indicated that the participants were not throwing with less force for fear of damaging the equipment, they were asked a fifth and final question:

- 5 The results from my experiment indicate that you were not throwing the ball with as much force as you might have in a real throwing exercise. Were you aware that you might be throwing with less force? (And if yes, do you know why?)

Responses are shown in table 5.10:

Table 5.10: Participant’s responses to the fifth follow-up questions asked to them.

ID	Response
1	I think that yes, I was aware I might throwing it with less force. The reason for that is that I didn’t have any feedback (even during training) but how far I was throwing the ball (I couldn’t see where the ball was going on the screen).
2	(no response given)
3	Yes I was aware. I think I was throwing the ball with enough force to reach the screen, and not at the target displayed on the screen. Perhaps in the future you could do a real throwing exercise just before the experiment.
4	When starting the experiments I felt involved when passing the ball and not acting, but it may have become a bit tedious/long-winded. Failing that I must be a bit of a wimp.
5	May be due to the tethering and not wanting to damage my wrist, plus as it wasn’t ‘real life’ then I may not have thought needed to throw using max strength.
6	Yes I was aware, I’d say the reason for that is because usually a lot of the force and distance comes from a combination of all of my upper body going into the throw and using that ‘twist’ to add the fastest spin too. In the simulator there was absolutely no upper body, it was all in the arms, and since the computer couldn’t really understand the fast spins. I ended up throwing it with virtually no spin at all which felt unnatural.

The participants provided some interesting responses. Three of the participants indicated that they were aware that they were throwing with less force, while the other two appeared to have been unaware and provided speculative responses.

Participant 1 noted that he believed it was due to lack of feedback, which was also an issue noted by participants during experiment 1.

Participant 3 said he believed he was throwing with “enough force to reach the screen”, but if he was not concerned with damaging the equipment it would appear he was not experiencing a level of immersion that would allow him to forget that the screen was there. He also suggested a real throwing exercise might have been beneficial.

Participant 6 mentioned that the system was failing to “understand the fast spins” in reference to the fact that he had to start again after a few trials as the system was unable to capture the data from his throws. He also indicated that the act of passing forward was making the task more difficult than if it had been more like the real task of rugby passing, where turning his upper body assisted with the pass.

Participant 4 stated he had “felt involved at the beginning but got tired”, which may also have affected the other participants, as they were required to do a long series of physical trials and only one regularly played rugby currently.

5.7 Summary

Experiment 2 confirmed that the throwing task is problematic for the users. Given that the participants responded well to distance judgements in the verbal task, it would appear that they are able to distinguish the distances acceptably, but that some factor of the throwing task is preventing the same performance. It is possible that the participants feel uncomfortable throwing the ball towards a screen, despite the use of the bungee, though further experiments would be required to determine this.

A Presence Questionnaire was administered to the participants after the second experiment was performed, which showed that most of the participants had a highly positive attitude towards their immersion experience.

Asking the participants some questions after the experiment revealed that the throwing task had been difficult and unnatural to some of them, but none of them indicated fear of damaging the equipment.

The results of the experiment will now be considered in the context of the research questions:

- What types of skills appear to be best suited to training in VEs?
- Which technologies achieve the best results in specific performance measures?

- Should stereoscopic displays be used?
- What does it cost?

As with experiment 1, the equipment used was very expensive, and even then the tracking systems struggled with the movement of the ball.

The use of stereoscopy was again found to not significantly affect the perceived target distances, while the shape of the screen was found to significantly affect estimated distance.

While the skill of passing may or may not be suited to training in a VE, the adapted task of throwing the ball at the targets was not well received as it was too different to reality for several of the participants.

Chapter 6

Discussions and Conclusion

In this final chapter, the research undertaken in this thesis is summarised, and the results of the two experiments are compared. The developed system is then discussed in the context of the eight research questions originally considered in chapter 2. From the issues raised during the two experiments and the discussion regarding the research questions, some potential future work is identified.

6.1 Summary

We have developed a prototype virtual environment for rugby skills training (VERST) in which users can practise their throwing skills by aiming a real ball at a virtual target. Three major versions of the system have been implemented (all programming was done by the author of this thesis with the exception of the Virtools portion of the system for experiment 1): the prototype version, programmed in XVR and VRPN, was used as a development platform and acted as a proof of concept for the idea. A collaborating team at CRVM in Aix-Marseille University implemented the second version using Virtools linked to a C++ implementation of the physics engine (written by the author of this thesis), for an experiment to determine whether users could correctly identify the virtual target distances, as evidence to the contrary already existed (Renner et al., 2014). The final version was implemented in Unity and MiddleVR, for a second experiment to follow up on the results of the first.

Two experiments were performed to measure the users' perception of egocentric distance to the virtual targets. The experiments investigated the following research questions for VERST:

- What types of skills appear to be best suited to training in VEs?
- Which technologies achieve the best results in specific performance measures?
- Should stereoscopic displays be used?
- What does it cost?

The first experiment, performed in a CAVE at CRVM, was based on four participants throwing the tethered ball at 280 targets under three experimental variables: the position of the participant relative to the screen, the use of the floor screen and the use of stereoscopy. None of the variables were found to significantly affect the distance judgements made by the participants, but the distances were almost all in the range of

6–9 metres, regardless of the distance of the targets (which were in a range of 5–15 metres) and not scaled to the target distances. This suggested that perhaps the participants were just throwing with the same force every time, and perhaps seeing no variation in target distance. The participants often threw in the wrong direction by a small amount, which may have been an effect of the shape of the screen – as they were facing diagonally across from the corner of the CAVE – or it may have been an effect of the physical distance from the participant to the screen.

To determine whether the shape of the screen had affected the distance or direction of the throws, the experiment was replicated using a flat screen. In this second experiment, the experimental variables used were the physical distance of the participant from the screen and the use of stereoscopy.

The second experiment contained three tasks: a throwing task (as in experiment 1), a verbal estimation task, and a presence questionnaire. The verbal experiment was performed using the same flat screen as the throwing experiment and a curved screen. For the second experiment, a smaller range of distances was used: 5–10 metres. The throwing task could not be performed using the curved screen due to the tracking system having a frame rate that was too low to capture the ball in flight. We intended to use the verbal task to attempt to draw parallels in the data from the flat screen to try to predict the users' behaviour in the curved screen for the throwing task.

The throwing task in experiment 2, despite the use of different participants and a different screen shape, yielded very similar results to the first. This time, however, the use of stereoscopy was found to produce results with a less compressed range of distances, suggesting the non-stereo targets all appeared to be at a similar distance. This data may have been anomalous due to the grouping, though, as it transpired the use of stereoscopy had no effect on the individual target distances. The position of the participant also proved to be significant, with the range of throws made from the medium position falling short of the target range.

In the verbal task of experiment 2, the estimates made by the participants were

scaled to the target distances. But while we had expected the participants to underestimate the distances based on previous work (Renner et al., 2014), the participants overestimated the distances. There proved to be a significant difference in the estimations between the curved and flat screens, with results from the curved screen generally overestimated by 2 metres and results from the flat screen generally overestimated by 3–4 metres.

For the verbal task, the use of stereoscopy was not found to be significant, but the position of the participant was highly significant. The amount of overestimation by the participants was strongly correlated with the physical distance between them and the screen. A Presence Questionnaire was administered to the participants after the second experiment was performed, which showed that most of the participants had a highly positive attitude towards their immersion experience.

6.2 Comparing the Results of Experiments 1 and 2

Two experiments were performed to measure the users' perception of egocentric distance to virtual targets. The first experiment, performed in a CAVE, was based on four participants throwing the tethered ball at 280 targets at 5–15 m under three experimental variables: the position of the participant relative to the screen, the use of the floor screen and the use of stereoscopy. The second experiment consisted of a throwing task performed using a flat screen, and a verbal estimation task performed using flat and curved screens. Experimental variables were the distance of the participant from the screen and the use of stereoscopy. For the second experiment, a smaller range of distances was used: 5–10 m.

The use of stereoscopy was not found to significantly affect the verbal estimations or thrown distances from either experiment. Neither was the use of the floor screen in experiment 1. The position of the participants did affect the distances perceived: in the first experiment, the position of the participant affected the direction of the throw,

while for the second experiment, the further the participant from the screen, the further they perceived targets.

The mean estimated distance data (throwing and verbal estimates) from both experiments are plotted in figure 6.1. That the participants were able to perceive a scaling in distance during the verbal experiment, but not to do the same during the throwing tasks suggested that they may have been uncomfortable with throwing the ball towards the screen. To attempt to learn more about the participants' reactions, they were asked several follow-up questions.




Figure 6.1: Comparison of the average estimated distance between the verbal and throwing tasks in the two experiments as a function of the intended distance. The lighter-coloured dashed line represents the 1:1 ratio of the two axes.

At odds with the hypothesis, the participants all claimed that they had no fear of damaging the equipment. Some of the participants did say they had found the task to be difficult or unnatural, mainly because it was not similar enough to the real task of passing in rugby, where the pass is thrown backwards whilst running.

It is likely, then, that the task was poorly designed for the skill it was attempting to train. If the VE had been training throwing skills out of the context of rugby, then the results of the throwing task may have been better. A participant also suggested that a proper warm-up task of throwing the ball at real-life targets may also have assisted. In

both experiments, participants stated that not seeing the virtual ball (and not receiving any feedback on the distance thrown) made the task harder and “strange”. The ball was not shown during the experiment to prevent ‘training to the simulator’, and so perhaps showing the ball would also improve results by making the system more natural to use.

6.3 Research Questions

In chapter 1, it was hypothesised that it would be possible to develop a VE that can effectively train ball passing skills in rugby. The two experiments performed using VERST indicated that it has not yet reached a stage in which the hypothesis can be accepted, but do indicate that that this could happen with further improvements (discussed in future work).

Returning to the eight research questions asked in chapter 2, we can now address these questions in the context of the VERST project (the two experiments focussed on the bolded questions):

1. Do sports skills reliably transfer from VE training conditions to real-world scenarios?
2. **What types of skills appear to be best suited to training in VEs?**
3. **Which technologies achieve the best results in specific performance measures?**
4. **Should stereoscopic displays be used?**
5. Is high fidelity always better?
6. How, what and when should feedback be delivered to the learner?
7. When is a VE not successful in teaching motor skills?
8. **What does it cost?**

6.3.1 Do Sports Skills Reliably Transfer from Virtual Environments Training Conditions to Real-World Scenarios?

It would appear that participants in the experiments did not successfully transfer their real-world skills into the virtual world, suggesting they are unable to disregard the knowledge that they are in a VE. While the participants did appear to exhibit some perceptual blindness (high scores on the Presence Questionnaire and one participant stating he had felt “involved” with the task), if they cannot perform the skill as they normally would, then they will be unable to work on training that skill in VERST.

6.3.2 What Types of Skills Appear to be Best Suited to Training in Virtual Environments?

While blind throwing of a bean bag has been previously used as distance perception task (Sahm et al., 2005), and was found to be accurate, throwing without any visual obstruction does not appear to have worked in the case of VERST. Two experiments involving different participants with different hardware produced very similar results, indicating there is an issue with the task to be performed. Previous research has found that blind throwing produces comparable results to blind walking (Sahm et al., 2005), and that timed imagined walking is comparable to verbal estimation (Piryankova et al., 2013), but the verbal and throwing results from both experiments in VERST were not at all comparable. While a small number of participants may be the cause of this, it could also be issues with the hardware or the physics engine that have not been found.

It is likely that the poor results from the throwing task are due to the task being changed too much during transfer from reality to virtuality. If the participants had been able to move and pass the ball backwards, perhaps VERST might have produced better results.

6.3.3 Which Technologies Achieve the Best Results in Specific Performance Measures?

Three different display technologies were compared during the experiments: a CAVE, a flat powerwall and a curved powerwall. Our attempt to compare the distance perception of participants across the three screens was inconclusive. The results in the throwing task were similar in the CAVE and flat screen hardware setups. Given that the verbal results in the curved screen were overestimated slightly less than those using the flat screen, it is suggested that the throwing results might have improved slightly in the curved screen setup, but this would require further work to verify.

Using a real ball as a haptic interaction device should have made the task of throwing more natural to the participants. There were two major issues associated with the use of the real ball: firstly, the tracking system had to have a sufficiently high capture rate, and secondly, the problem of tethering the ball to prevent any equipment damage. 55–60 Hz was not found to be an acceptable capture rate from the PST (55 Hz) and the ART system in the curved screen laboratory (60 Hz), whereas the Vicon system capturing at 120–250 Hz provided enough data. Though it should be noted that participant 6 had difficulty with the Vicon system being unable to sufficiently track the ball when he was using spin, and so it is possible even the tracking systems with higher capture rates would not be acceptable.

Three methods of securing the ball were used: the first, for the prototype system, tethered the ball to a legless table top with a bungee; the second used climbing rope to secure it to the ceiling rafters of the building; the third involved a wrist strap and a bungee worn by the user. The table top provided a secure base, and the bungee pulled the ball downwards and away from the participant, however it was very large and not very portable for the experiments which were conducted far from Bangor University. The ceiling tether was problematic as the ball would swing back heavily towards the participant, and would often distract them from the VE. It also tangled with the retrore-

flective marker rig frequently. The wrist strap and bungee were difficult for the user to grow accustomed to, as the ball would snap back towards them when the bungee flexed. This was also very off-putting to the users, but very portable. If VERST were to be more permanently installed into a room, fixing the ball to the floor seems to be the best option.

6.3.4 Should Stereoscopic Displays be Used?

In both experiments performed using VERST, the use of stereoscopy was not found to be a significant factor in the estimation of depth, which is consistent with previous findings using similar egocentric distance perception experiments (Renner et al., 2014; Roumes et al., 2001; Eggleston et al., 1996; Willemsen et al., 2008). This may have been due to the scene lacking virtual objects close to the participant, and therefore not providing enough depth cues to make stereoscopic presentation useful to their estimations.

If the use of stereoscopy had been found to increase the users' sense of presence, it would likely contribute to a better training environment. This could be addressed in future work.

6.3.5 Is High Fidelity Always Better?

In terms of perceptual fidelity, as with the use of stereoscopy, previous work has neither confirmed nor denied that higher quality graphics influence the perception of distance in a VE. In the case of VERST, the original XVR testbed environment was abandoned in favour of one created in Unity with higher quality of graphics. We believe that the higher the quality of the graphics in the environment, the more the user will feel immersed. The greater sense of reality and presence can only improve the quality of VERST as a training tool. The lower-quality version was not used in any experiments, so it is only possible to speculate on this matter.

The functional fidelity of VERST was not fully explored, as the experiments focussed on the ability of the user to perceive the targets at the intended distance. VERST included motion parallax, which allowed users to walk around in front of the screen and view the scene with accurate geometric projection. Throwing the ball was the only other interaction available to the user, as it was the only necessary interaction for the basic task of throwing at the targets.

6.3.6 How, What and When Should Feedback be Delivered to the Learner?

Feedback on the outcome of the trials was not used with VERST. Participants in the initial experiment at CRVM did state that they found it “strange” to not be able to see the virtual ball, but it is possible that they missed it because they were expecting it. In the second experiment, the participants were not informed that there was a virtual ball that would be hidden from them, and none of them commented on the lack of it during the experiment, though one did suggest in response to the follow-up questions that this may have been why the task was difficult for him. The final participant was interested and asked further questions about VERST after he had completed his trials and was shown the version with the ball; conversely, this was the version he found “strange”.

We had expected the pitch marking to be used by participants in both of the experiments, as a general environmental cue, and had expected a high rate of correct results, particularly in the verbal task. Two of the participants in the Hartree Centre experiments questioned whether they were “allowed” to use them, with one stating that he had to as he could only estimate distances by points of reference. The others may have used the 10 metre line (the most visible from their location) as a point of reference, but the estimations were still mainly incorrect. This may be associated with work by Armbrüster et al. (2008) which showed that placing a virtual tape measure in the environment did not aid the subjects’ perceptions of egocentric distance.

6.3.7 When is a Virtual Environment Not Successful in Teaching Motor Skills?

As with the first question, if the skill cannot be transferred from reality to virtuality, then the skill cannot be trained. In the case of VERST, as we have been unable to reach that stage, it cannot be considered to successfully teach a motor skill.

6.3.8 What Does it Cost?

Costly equipment and software proved to be problematic, but fortunately for this project, research funding and opportunities were available to use. While the initial version of VERST was developed using free software and a relatively cheap tracking device, the two experimental models required very expensive hardware – a CAVE, Vicon and ART tracking system, curved and flat rear-projected powerwalls – and software – Virtools to produce. While Unity and MiddleVR, the software for the third version were free, the stereo features could only be accessed using Unity Pro, and the free version of MiddleVR is extremely restrictive, resulting in the use of trial licenses. While it would have been possible to build VERST using free software and cheap hardware, it would have been of much lower quality.

6.3.9 Conclusion

Many of the responses to the research questions in the context of VERST are negative, but we believe that the research undertaken to build VERST will provide an important platform for future work. To our knowledge, there are no other VEs that have attempted to train rugby passing skills, and no other sports training VEs have been found which allow the user to physically throw an object to interact with the VE. Chong and Croft's lineout training environment used video footage as opposed to a virtual scene, and the user threw at a projected stone wall rather than expensive equipment (Chong and Croft, 2009).

The hypothesis “A virtual environment can be used to effectively train ball passing skills in rugby”, can not be accepted without further evaluation of VERST which is beyond the scope of this research. While it may be possible in future to fully-develop a VE to train rugby skills, VERST has not yet fulfilled the hypothesis.

6.4 Future Work

6.4.1 Further Experiments

An experiment to examine the users’ reactions to differently configured versions of the physics engine would provide insight into the visual accuracy of the predicted flight paths. It would be possible to place the subject in VERST and present them with a number of calculated flight paths, each calculated using different drag and lift coefficients, and to ask the participant to report which appeared most realistic.

Measuring the accuracy of a subject in the real world before and after immersion into VERST for training would provide an indication of any negative transfer, and to also attempt to validate the system. This would require each subject to be observed – and possibly tracked – in great detail to determine the accuracy of throws in the real world. The target distances would need to be matched in reality and virtuality, and the real targets moved between each trial, making this a time-consuming experiment to perform.

To further investigate the cause of the poor results from the throwing task, and to investigate if the task of throwing forwards (as a move which is illegal in a game of rugby) was the cause of the difficulties experienced by the participants, it is possible that the system may be adapted to allow users to throw the ball sideways to the distance of the target that they see. Although this task was considered to form part of the experiment at the Hartree Centre, there was not enough time to implement it. This may form a more natural task for the rugby players, as they must throw sideways and behind their current position during play.

It may also be possible to have the participants facing away from the screen and turning their upper body to throw towards the target in the virtual scene behind them.

6.4.2 Improvements to the Perceptual and Functional Fidelity

There is a great deal of further work that could be done with VERST. While the quality of the graphics in the system is good, it could be further improved with time. Adding sounds such as crowd noises would likely improve the perceptual fidelity of the system further. No sounds were used in either the prototype or the final version, and were only used in the second version to mark the beginning and failure of trials. Little work was done on the functional fidelity of the system, there was no user interface in the final version, and the system should ideally be installed in a fixed location, to allow for proper development with the hardware. The tracking systems, in particular, should be worked on to improve the capture of the ball. The user should also be able to receive some form of feedback, during or post-immersion.

6.4.3 Improvements to the Physics Engine

The physics engine should undergo further testing. While we lacked the facilities to perform a full analysis of the model and check the integrity, it forms the core of the system, and should be thoroughly verified. While we had intended to implement weather effects into VERST, it was not possible in the time scale of the thesis.

6.4.4 Implementation of Scenarios and Training Activities

Another feature not implemented in the time scale of the thesis was the introduction of scenarios and training activities (other than the simple throwing forward drill). Scenarios could include the addition of virtual players who would act as targets, both stationary and moving. Multiple targets could exist in the scene, as well as different

styles of throwing exercises, such as throwing over a target or through a hole in a target (similar to the Brian Moore's Mouth Challenge).

6.4.5 Attempting to Solve the Poor Results from the Throwing Task

While HMDs were discounted at the beginning of the project, it would be of interest to perform the throwing experiment using a lightweight HMD (such as the Oculus Rift) and investigate the suitability of the throwing task in VEs. If our hypothesis that the users may be fearful of hitting the screen is correct, performing the throwing task in a spacious room using the HMD should allow the users to throw at full strength.

Bibliography

- Ahn, S. H., Kang, K. B., Kim, E. J., Kim, S.-J., Lee, J.-W., and Song, C. G. (2006). Interactive 3D golf simulator. *ICHIT-AMTA, November*, pages 10–11.
- Akeley, K., Watt, S. J., Girshick, A. R., and Banks, M. S. (2004). A stereo display prototype with multiple focal distances. In *ACM SIGGRAPH 2004*, volume 1, page 804, New York, New York, USA. ACM Press.
- Alahakone, A. and Senanayake, S. A. (2009). Vibrotactile feedback systems: Current trends in rehabilitation, sports and information display. In *IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM2009)*, pages 1148–1153. IEEE.
- Alam, F., Chowdhury, H., Moria, H., and Fuss, F. K. (2010). A comparative study of football aerodynamics. *Procedia Engineering*, 2(2):2443–2448.
- Alam, F., Smith, S., Djamovski, V., Yang, J., Kumar, A., Farabi, R., Ehsan, S., and Mainuddin, G. (2011). Aerodynamics of american footballs under crosswinds. In *International Conference on Mechanical Engineering and Renewable Energy 2011*, pages 1–6. Chittagong University of Engineering and Technology.
- Alam, F., Subic, A., Watkins, S., Naser, J., and Rasul, M. (2008). An experimental and computational study of aerodynamic properties of rugby balls. *WSEAS Transactions on Fluid Mechanics*, 3(3):279–286.

- Alam, F., Subic, A., Watkins, S., and Smits, A. J. (2009). Aerodynamics of an australian rules foot ball and rugby ball. In *Computational Fluid Dynamics for Sport Simulation*, pages 103–127. Springer.
- Alam, F., Watkins, S., and Subic, A. (2003). Aerodynamics of a rugby ball and australian rules foot ball. In *International Conference on Mechanical Engineering (ICME2003)*, pages 26–28, Dhaka, Bangladesh.
- Alaways, L. W. (1998). *Aerodynamics of the curve-ball: an investigation of the effects of angular velocity on baseball trajectories*. PhD thesis, University of California.
- American Educational Research Association, American Psychological Association, and National Council on Measurement in Education and Joint Committee on Standards for Educational and Psychological Testing (US) (1999). *Standards for educational and psychological testing*. Amer Educational Research Assn.
- Armbrüster, C., Wolter, M., Kuhlen, T., Spijkers, W., and Fimm, B. (2008). Depth perception in virtual reality: distance estimations in peri-and extrapersonal space. *Cyberpsychology & Behavior*, 11(1):9–15.
- Arora, S., Aggarwal, R., Sevdalis, N., Moran, A., Sirimanna, P., Kneebone, R., and Darzi, A. (2010). Development and validation of mental practice as a training strategy for laparoscopic surgery. *Surgical endoscopy*, 24(1):179–187.
- Balci, O. (1994). Validation, verification, and testing techniques throughout the life cycle of a simulation study. *Annals of operations research*, 53(1):121–173.
- Balci, O. (1995). Principles and techniques of simulation validation, verification, and testing. In *Simulation Conference Proceedings, 1995. Winter*, pages 147–154. IEEE.
- Balci, O. (1997). Verification validation and accreditation of simulation models. In *Proceedings of the 29th conference on Winter simulation*, pages 135–141. IEEE Computer Society.

- Banks, M. S., Read, J. C., Allison, R. S., and Watt, S. J. (2012). Stereoscopy and the human visual system. *SMPTE motion imaging journal*, 121(4):24–43.
- Barbati, A. (2014). Archery target texture. https://commons.wikimedia.org/wiki/File:Archery_Target_80cm.svg, Date Accessed: 02/01/14.
- Barber, S., Chin, S., and Carré, M. (2009). Sports ball aerodynamics: a numerical study of the erratic motion of soccer balls. *Computers & Fluids*, 38(6):1091–1100.
- Barber, S., Haake, S. J., and Carré, M. (2006). Using CFD to understand the effects of seam geometry on soccer ball aerodynamics. In *The Engineering of Sport 6*, pages 127–132. Springer.
- Battig, W. F. (1966). Facilitation and interference. *Acquisition of skill*, pages 215–244.
- Battig, W. F. (1979). The flexibility of human memory. *Levels of processing in human memory*, pages 23–44.
- BBC Wales Scrum V (2014). Brian Moore texture source image. http://news.bbcimg.co.uk/media/images/51758000/jpg/_51758119_mooremouth_grab.jpg, Date Accessed: 02/01/14.
- Bideau, B., Kulpa, R., Ménardais, S., Fradet, L., Multon, F., Delamarche, P., and Arnaldi, B. (2003). Real handball goalkeeper vs. virtual handball thrower. *Presence: teleoperators and virtual environments*, 12(4):411–421.
- Bideau, B., Kulpa, R., Vignais, N., Brault, S., Multon, F., and Craig, C. (2010a). Using virtual reality to analyze sports performance. *Computer Graphics and Applications, IEEE*, 30(2):14–21.
- Bideau, B., Kulpa, R., Vignais, N., Brault, S., Multon, F., and Craig, C. (2010b). Virtual reality, a serious game for understanding performance and training players in sport. *IEEE Computer Graphic Applications*, 30(2):14–21.

- Bideau, B., Multon, F., Kulpa, R., Fradet, L., and Arnaldi, B. (2004a). Virtual reality applied to sports: do handball goalkeepers react realistically to simulated synthetic opponents? In *Proceedings of the 2004 ACM SIGGRAPH international conference on Virtual Reality continuum and its applications in industry*, pages 210–216. ACM.
- Bideau, B., Multon, F., Kulpa, R., Fradet, L., Arnaldi, B., and Delamarche, P. (2004b). Using virtual reality to analyze links between handball thrower kinematics and goalkeeper's reactions. *Neuroscience letters*, 372(1):119–122.
- Bierbaum, A., Just, C., Hartling, P., Meinert, K., Baker, A., and Cruz-Neira, C. (2001). VR Juggler: A virtual platform for virtual reality application development. In *Virtual Reality, 2001. Proceedings. IEEE*, pages 89–96. IEEE.
- Boehm, B. W. (1979). Software engineering: R&D trends and defense needs. *Research directions in software technology*, 1:977.
- Bolte, B., Zeidler, F., Bruder, G., Steinicke, F., Hinrichs, K., Fischer, L., and Schorer, J. (2010). A virtual reality handball goalkeeper analysis system. In *Proceedings of the Joint Virtual Reality Conference of EuroVR-EGVE-VEC (JVRC'10). Stuttgart, Germany*, pages 1–2.
- Brady, F. (2004). Contextual interference: A meta-analytic study 1. *Perceptual and motor skills*, 99(1):116–126.
- Braut, S., Bideau, B., Craig, C., and Kulpa, R. (2010). Balancing deceit and disguise: How to successfully fool the defender in a 1 vs. 1 situation in rugby. *Human movement science*, 29(3):412–425.
- Braut, S., Bideau, B., Kulpa, R., and Craig, C. (2009). Detecting deceptive movement in 1 vs. 1 based on global body displacement of a rugby player. *International Journal of Virtual Reality*, 8(4):31–36.

- Bray, K. and Kerwin, D. G. (2003). Modelling the flight of a soccer ball in a direct free kick. *Journal of sports sciences*, 21(2):75–85.
- Bridgeman, B., Gemmer, A., Forsman, T., and Huemer, V. (2000). Processing spatial information in the sensorimotor branch of the visual system. *Vision Research*, 40(25):3539–3552.
- Brook, P., Croft, H., and Mann, S. (2007). Laser based line-out simulator. In *Proceedings of the 20th Annual Conference of the National Advisory Committee on Computing Qualification (NACCQ'07)*. Nelson, National Advisory Committee on Computing Qualifications, New Zealand, page 265.
- Brunnett, G., Rusdorf, S., and Lorenz, M. (2006). V-Pong: an immersive table tennis simulation. *Computer Graphics and Applications, IEEE*, 26(4):10–13.
- Burdea, G. and Coiffet, P. (2003). Virtual reality technology. *Presence: Teleoperators and virtual environments*, 12(6):663–664.
- Burglund, B. (2011). Golf ball flight dynamics. pages 1–19.
- Carrozzino, M., Tecchia, F., Bacinelli, S., Cappelletti, C., and Bergamasco, M. (2005). Lowering the development time of multimodal interactive application: the real-life experience of the XVR project. In *Proceedings of the 2005 ACM SIGCHI International Conference on Advances in computer entertainment technology*, pages 270–273. ACM.
- Cater, K., Chalmers, A., and Ledda, P. (2002). Selective quality rendering by exploiting human inattention blindness: looking but not seeing. In *Proceedings of the ACM symposium on Virtual reality software and technology*, pages 17–24. ACM.
- chai3D.org (2014). CHAI 3D. <http://www.chai3d.org/>, Date Accessed: 21/11/11.

- Chiaet, J. (2013). Getting on the ball: How the FIFA 14 soccer video game finally got its physics right. <http://www.scientificamerican.com/article.cfm?id=getting-on-the-ball-how-soccer-video-game-got-physics-right>, Date Accessed: 05/12/13.
- Chiviacosky, S. and Wulf, G. (2007). Feedback after good trials enhances learning. *Research Quarterly for Exercise and Sport*, 78(2):40–47.
- Chong, A. K. and Croft, H. G. (2009). A photogrammetric application in virtual sport training. *The Photogrammetric Record*, 24(125):51–65.
- Chong, A. K., Milburn, P., Newsham West, R., ter Voert, M., and Croft, H. (2008). Recent practical applications of close-range photogrammetry for complex motion study. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXVII(Part B5):921–926.
- Chua, P. T., Crivella, R., Daly, B., Hu, N., Schaaf, R., Ventura, D., Camill, T., Hodgins, J., and Pausch, R. (2002). Mastermotion: Full body wireless virtual reality for tai chi. In *ACM SIGGRAPH 2002 conference abstracts and applications*, pages 214–214. ACM.
- Chua, P. T., Crivella, R., Daly, B., Hu, N., Schaaf, R., Ventura, D., Camill, T., Hodgins, J., and Pausch, R. (2003). Training for physical tasks in virtual environments: Tai Chi. In *Virtual Reality, 2003. Proceedings. IEEE*, pages 87–94. IEEE.
- Chung, J., Xu, K., Colaco, A., Schmandt, C., and Li, V. O. (2010). My second bike: a TV-enabled social and interactive riding experience. In *Consumer Communications and Networking Conference (CCNC), 2010 7th IEEE*, pages 1–5. IEEE.
- Cooke, A. J. (2000). An overview of tennis ball aerodynamics. *Sports engineering*, 3(2):123–129.

- Craig, C. M., Berton, E., Rao, G., Fernandez, L., and Bootsma, R. J. (2006). Judging where a ball will go: the case of curved free kicks in football. *Naturwissenschaften*, 93(2):97–101.
- Creem-Regehr, S. H., Willemsen, P., Gooch, A. A., and Thompson, W. B. (2005). The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual environments. *Perception*, 34(2):191–204.
- Cross, R. (1987). *Ball Trajectories: Factors Influencing the Flight of the Ball*.
- Djamovski, V. and Alam, F. (2011). An experimental aerodynamic study of a series of rugby balls. In *International Conference on Mechanical Engineering*, pages 1–6. BUET.
- Djamovski, V., Rosette, P., Chowdhury, H., Alam, F., and Steiner, T. (2012). A comparative study of rugby ball aerodynamics. *Procedia Engineering*, 34:74–79.
- EA Sports (2014a). FIFA Player Impact Engine. <http://www.ea.com/news/fifa12-player-impact-engine-01>, Date Accessed: 06/12/13.
- EA Sports (2014b). NHL®14 Collision Physics. <http://www.easports.com/nhl/news-updates-gameplay/article/nhl14-developer-diary-collision-physics>, Date Accessed: 06/12/13.
- Eaves, D. L., Breslin, G., Van Schaik, P., Robinson, E., and Spears, I. R. (2011). The short-term effects of real-time virtual reality feedback on motor learning in dance. *Presence: Teleoperators and Virtual Environments*, 20(1):62–77.
- Eggleston, R. G., Janson, W. P., and Aldrich, K. A. (1996). Virtual reality system effects on size-distance judgements in a virtual environment. In *Virtual Reality Annual International Symposium, 1996., Proceedings of the IEEE 1996*, pages 139–146. IEEE.

- Elliott, D., Chua, R., Pollock, B. J., and Lyons, J. (1995). Optimizing the use of vision in manual aiming: The role of practice. *The Quarterly Journal of Experimental Psychology*, 48(1):72–83.
- Emoto, M., Niida, T., and Okano, F. (2005). Repeated vergence adaptation causes the decline of visual functions in watching stereoscopic television. *Display Technology, Journal of*, 1(2):328–340.
- Ericsson, K. A., Krampe, R. T., and Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological review*, 100(3):363.
- Essilor (2014). Digital corneal reflection pupilometer. <http://www.essilor.co.uk/Instruments/DispensingAids/Pages/Pupilometer.aspx>, Date Accessed: 02/01/14.
- Federal Aviation Administration, U.S. Department of Transportation (2014). Specification for Wind Cone Assemblies. http://www.faa.gov/documentLibrary/media/advisory_circular/150-5345-27D/150_5345_27d.pdf, Date Accessed: 02/01/14.
- Fels, S., Kinoshita, Y., Chen, T.-P. G., Takama, Y., Yohanan, S., Gadd, A., Takahashi, S., and Funahashi, K. (2005). Swimming across the pacific: a VR swimming interface. *Computer Graphics and Applications, IEEE*, 25(1):24–31.
- Fincham, E. and Walton, J. (1957). The reciprocal actions of accommodation and convergence. *The Journal of Physiology*, 137(3):488–508.
- Fink, P. W., Foo, P. S., and Warren, W. H. (2009). Catching fly balls in virtual reality: A critical test of the outfielder problem. *Journal of Vision*, 9(13).
- Fofi, D., Sliwa, T., and Voisin, Y. (2004). A comparative survey on invisible structured light. In *Proceedings of SPIE*, volume 5303, pages 90–98.

- Fuss, F. K. and Troynikov, O. (2012). Anisotropic friction of rugby ball surfaces. *Proceedings of the Institution of Mechanical Engineers, Part J: Journal of Engineering Tribology*, 226(7):608–615.
- Giant Army (2014). Universe Sandbox. <http://universesandbox.com/>, Date Accessed: 06/12/13.
- Göbel, S., Geiger, C., Heinze, C., and Marinos, D. (2010). Creating a virtual archery experience. In *Proceedings of the International Conference on Advanced Visual Interfaces*, pages 337–340. ACM.
- Goff, J. E. and Carré, M. J. (2010). Soccer ball lift coefficients via trajectory analysis. *European Journal of Physics*, 31(4):775–784.
- Goodale, M. A. and Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in neurosciences*, 15(1):20–25.
- Gourishankar, V., Srimathveeravalli, G., and Kesavadas, T. (2007). Hapstick: A high fidelity haptic simulation for billiards. In *EuroHaptics Conference, 2007 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2007. Second Joint*, pages 494–500. IEEE.
- Grantcharov, T. P., Kristiansen, V., Bendix, J., Bardram, L., Rosenberg, J., and Funch-Jensen, P. (2004). Randomized clinical trial of virtual reality simulation for laparoscopic skills training. *British Journal of Surgery*, 91(2):146–150.
- Gray, R. (2002). Behavior of college baseball players in a virtual batting task. *Journal of Experimental Psychology: Human Perception and Performance*, 28(5):1131.
- Griffiths, I., Evans, C., and Griffiths, N. (2005). Tracking the flight of a spinning football in three dimensions. *Measurement Science and Technology*, 16(10):2056–2065.

- Groen, A. (2013). NASAs Jet Propulsion Lab is obsessed with a certain game, and I bet you can guess what it is. <http://penny-arcade.com/report/article/nasas-jet-propulsion-lab-is-obsessed-with-a-certain-game-and-i-bet-you-can->, Date Accessed: 05/12/13.
- Hays, R. T., Jacobs, J. W., Prince, C., and Salas, E. (1992). Flight simulator training effectiveness: A meta-analysis. *Military Psychology*, 4(2):63–74.
- Henry, F. M. (1968). Specificity vs. generality in learning motor skill. In Brown, R. and Kenyon, G. S., editors, *Classical studies on physical activity*, pages 331–40. Prentice Hall, Englewood Cliffs, NJ.
- Hillis, J. M., Watt, S. J., Landy, M. S., and Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision*, 4(12).
- Hoffman, D. M., Girshick, A. R., Akeley, K., and Banks, M. S. (2008). Vergence-accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of vision*, 8(3).
- Hoffman, D. M., Karasev, V. I., and Banks, M. S. (2011). Temporal presentation protocols in stereoscopic displays: flicker visibility, perceived motion, and perceived depth. *Journal of the Society for Information Display*, 19(3):271–297.
- Hoiem, D., Efros, A. A., and Hebert, M. (2005). Automatic photo pop-up. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 577–584. ACM.
- Howard, I. P. (2002). Seeing in depth, vol. 1: Basic mechanisms.
- Howard, I. P. and Rogers, B. J. (2002). *Seeing in depth, Vol. 2: Depth perception*. I Porteous Publishing, Toronto.
- Huegel, J. C. and O'Malley, M. K. (2010). Progressive haptic and visual guidance for training in a virtual dynamic task. In *Haptics Symposium, 2010 IEEE*, pages 343–350. IEEE.

- Hughes, M. and Franks, I. (2004). *Notational Analysis of Sport: Systems for Better Coaching and Performance in Sport*. Routledge, New York, 2nd ed edition.
- Interrante, V., Ries, B., and Anderson, L. (2006). Distance perception in immersive virtual environments, revisited. In *Virtual Reality Conference, 2006*, pages 3–10. IEEE.
- Jacobs, M. (2008). *Design and recognition of tactile feedback patterns for snowboarding*. PhD thesis, RWTH Aachen University.
- Jacobs, R. A. (1999). Optimal integration of texture and motion cues to depth. *Vision research*, 39(21):3621–3629.
- John, N. W. (2008). Design and implementation of medical training simulators. *Virtual reality*, 12(4):269–279.
- Kane, D., Held, R. T., and Banks, M. S. (2012). Visual discomfort with stereo 3d displays when the head is not upright. *Proceedings of Stereoscopic Displays and Applications XXIII, Burlingame, CA, USA*, page 828814.
- Keefe, B. D., Hibbard, P. B., and Watt, S. J. (2011). Depth-cue integration in grasp programming: no evidence for a binocular specialism. *Neuropsychologia*, 49(5):1246–1257.
- Kelly, P., Healy, A., Moran, K., and O'Connor, N. E. (2010). A virtual coaching environment for improving golf swing technique. In *Proceedings of the 2010 ACM workshop on Surreal media and virtual cloning*, pages 51–56. ACM.
- Khan, M. A. and Franks, I. M. (2000). The effect of practice on component submovements is dependent on the availability of visual feedback. *Journal of motor behavior*, 32(3):227–240.
- Khan, M. A., Franks, I. M., and Goodman, D. (1998). The effect of practice on the control of rapid aiming movements: Evidence for an interdependency between pro-

- programming and feedback processing. *The Quarterly Journal of Experimental Psychology: Section A*, 51(2):425–443.
- Kleijnen, J. P. (1995). Verification and validation of simulation models. *European Journal of Operational Research*, 82(1):145–162.
- Knepell, P. L. and Arangno, D. C. (1993). *Simulation validation: a confidence assessment methodology*, volume 15. John Wiley & Sons.
- Knill, D. C. (2005). Reaching for visual cues to depth: The brain combines depth cues differently for motor control and perception. *Journal of Vision*, 5(2).
- Knill, D. C. and Saunders, J. A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision research*, 43(24):2539–2558.
- Knörlein, B., Székely, G., and Harders, M. (2007). Visuo-haptic collaborative augmented reality ping-pong. In *Proceedings of the international conference on Advances in computer entertainment technology*, pages 91–94. ACM.
- Kotaku (2012). Kotaku Review: Kinect Sports. <http://kotaku.com/5679411/review-kinect-sports>, Date Accessed: 17/01/12.
- Lages, M. (2006). Bayesian models of binocular 3-d motion perception. *Journal of vision*, 6(4).
- Land, M. F. and McLeod, P. (2000). From eye movements to actions: how batsmen hit the ball. *Nature neuroscience*, 3(12):1340–1345.
- Law, A. M., Kelton, W. D., and Kelton, W. D. (1991). *Simulation modeling and analysis*, volume 2. McGraw-Hill New York.
- Lawrence, G. and Kingston, K. (2008). Skill acquisition for coaches. *An introduction to sports coaching: from science and theory to practice*, pages 16–27.

- Lee, T. D. and Magill, R. A. (1985). Can forgetting facilitate skill acquisition? *Differing perspectives in motor learning, memory, and control*, 27:3–22.
- Li, Y., Shark, L., Hobbs, S. J., and Ingham, J. (2010). Real-time immersive table tennis game for two players with motion tracking. In *Information Visualisation (IV), 2010 14th International Conference*, pages 500–505. IEEE.
- Lieberman, J. and Breazeal, C. (2007). TIKL: Development of a wearable vibrotactile feedback suit for improved human motor learning. *IEEE Transactions on Robotics*, 23(5):919–926.
- Liu, S., Cheng, D., and Hua, H. (2008). An optical see-through head mounted display with addressable focal planes. In *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*, pages 33–42.
- López-Moliner, J., Field, D. T., and Wann, J. P. (2007). Interceptive timing: Prior knowledge matters. *Journal of Vision*, 7(13).
- Love, G. D., Hoffman, D. M., Hands, P. J., Gao, J., Kirby, A. K., and Banks, M. S. (2009). High-speed switchable lens enables the development of a volumetric stereoscopic display. *Optics express*, 17(18):15716.
- MacKenzie, K. J., Hoffman, D. M., and Watt, S. J. (2010). Accommodation to multiple-focal-plane displays: Implications for improving stereoscopic displays and for accommodation control. *Journal of Vision*, 10(8).
- Mackrout, I. and Proteau, L. (2007). Specificity of practice results from differences in movement planning strategies. *Experimental Brain Research*, 183(2):181–193.
- Martens, T. and Ogle, K. (1959). Observations on accommodative convergence; especially its nonlinear relationships. *American journal of ophthalmology*, 47(1 Part 2):455.

- Matsuoka, Y. and Miller, L. C. (1999). Domestic rehabilitation and learning of task-specific movements. In *Sixth International Conference on Rehabilitation Robotics*, pages 177–182.
- Matsuura, H., Abe, N., Tanaka, K., and Taki, H. (2006). Constructing virtual air hockey game through the network. In *Advanced Information Networking and Applications, 2006. AINA 2006. 20th International Conference on*, volume 2, page 5. IEEE.
- Mechdyne Corporation (2014). CAVELib. <http://www.mechdyne.com/cavelib.aspx>, Date Accessed: 17/01/12.
- Mehta, R. and Pallis, J. (2001). The aerodynamics of a tennis ball. *Sports Engineering*, 4(4):177–189.
- Mehta, R. D. (1985). Aerodynamics of sports balls. *Annual Review of Fluid Mechanics*, 17(1):151–189.
- Mehta, R. D. and Koga, D. J. (2000). Cricket ball aerodynamics: myth versus science.
- Messick, S. (1995). Validity of psychological assessment: validation of inferences from persons' responses and performances as scientific inquiry into score meaning. *American Psychologist*, 50(9):741.
- Mestre, D. R., Dagonneau, V., Mercier, C. C.-S., and Maiano, C. (2011). Does virtual reality enhance exercise performance, enjoyment, and dissociation? An exploratory study on a stationary bike apparatus. *Presence: Teleoperators and Virtual Environments*, 20(1):1–14.
- Model created by user 'dizzyHARSH' (2014). Millennium Stadium model. <http://sketchup.google.com/3dwarehouse/details?mid=8267d16d1b1a4e24cca527397568d9ea>, Date Accessed: 02/01/14.
- Moeslund, T. B. and Granum, E. (2001). A survey of computer vision-based human motion capture. *Computer Vision and Image Understanding*, 81(3):231–268.

- Mori, M. (1970). The uncanny valley. *Energy*, 7(4):33–35.
- Napieralski, P. E., Altenhoff, B. M., Bertrand, J. W., Long, L. O., Babu, S. V., Pagano, C. C., Kern, J., and Davis, T. A. (2011). Near-field distance perception in real and virtual environments using both verbal and action responses. *ACM Transactions on Applied Perception*, 8(3):1–19.
- NASA Glenn Research Center (2014a). Boundary Layer. <http://www.grc.nasa.gov/WWW/k-12/airplane/boundlay.html>, Date Accessed: 02/01/14.
- NASA Glenn Research Center (2014b). The Drag Equation. <http://www.grc.nasa.gov/WWW/k-12/airplane/drageq.html>, Date Accessed: 02/01/14.
- NASA Glenn Research Center (2014c). The Lift Equation. <http://www.grc.nasa.gov/WWW/k-12/airplane/lifteq.html>, Date Accessed: 02/01/14.
- NASA Glenn Research Center (2014d). The Weight Equation. <http://www.grc.nasa.gov/WWW/k-12/airplane/wteq.html>, Date Accessed: 02/01/14.
- Nintendo (2014). Nintendo Wii Remote and Wii MotionPlus. <http://www.nintendo.com/wii/console/controllers>, Date Accessed: 18/01/12.
- Parks, T. E. (2012). Visual-illusion distance paradoxes: A resolution. *Attention, Perception, & Psychophysics*, 74(8):1568–1569.
- Perez, C. R., Meira Jr, C. M., and Tani, G. (2005). Does the contextual interference effect last over extended transfer trials? *Perceptual and motor skills*, 100(1):58–60.
- Pioch, N. J. (1995). Officer of the deck: validation and verification of a virtual environment for training.
- Piryankova, I. V., de la Rosa, S., Kloos, U., Bühlhoff, H. H., and Mohler, B. J. (2013). Egocentric distance perception in large screen immersive displays. *Displays*, 34(2):153–164.

- Potel, M. (1998). Motion sick in cyberspace. *Computer Graphics and Applications, IEEE*, 18(1):16–21.
- Proteau, L. and Cournoyer, J. (1990). Vision of the stylus in a manual aiming task: The effects of practice. *The Quarterly Journal of Experimental Psychology*, 42(4):811–828.
- Proteau, L. and Marteniuk, R. G. (1993). Static visual information and the learning and control of a manual aiming movement. *Human Movement Science*, 12(5):515–536.
- Proteau, L., Marteniuk, R. G., Girouard, Y., and Dugas, C. (1987). On the type of information used to control and learn an aiming movement after moderate and extensive training. *Human Movement Science*, 6(2):181–199.
- Push Square (2014). Image of Rugby World Cup 2011. <http://images.pushsquare.com/screenshots/39311/large.jpg>, Date Accessed: 02/01/14.
- Rae, W. J. (2003). Flight dynamics of an american football in a forward pass. *Sports Engineering*, 6(3):149–163.
- Ramic-Brkic, B. and Chalmers, A. (2010). Virtual smell: Authentic smell diffusion in virtual environments. In *Proceedings of the 7th International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa, AFRI-GRAPH '10*, pages 45–52, New York, NY, USA. ACM.
- Regan, D. and Gray, R. (2000). Visually guided collision avoidance and collision achievement. *Trends in Cognitive Sciences*, 4(3):99–107.
- Renner, R., Velichkovsky, B., and Helmert, J. (2014). The perception of egocentric distances in Virtual Environments - A Review. *ACM Computing Surveys*, To Appear.

- Roumes, C., Meehan, J. W., Plantier, J., and Menu, J.-P. (2001). Distance estimation in a 3-d imaging display. *The International Journal of Aviation Psychology*, 11(4):381–396.
- Ruffaldi, E., Filippeschi, A., Avizzano, C. A., Bardy, B., Gopher, D., and Bergamasco, M. (2011). Feedback, affordances, and accelerators for training sports in virtual environments. *Presence: Teleoperators and Virtual Environments*, 20(1):33–46.
- Rusdorf, S., Brunnett, G., Lorenz, M., and Winkler, T. (2007). Real-time interaction with a humanoid avatar in an immersive table tennis simulation. *Visualization and Computer Graphics, IEEE Transactions on*, 13(1):15–25.
- Rushton, S. K. and Wann, J. P. (1999). Weighted combination of size and disparity: a computational model for timing a ball catch. *Nature neuroscience*, 2(2):186–190.
- Sahm, C. S., Creem-Regehr, S. H., Thompson, W. B., and Willemsen, P. (2005). Throwing versus walking as indicators of distance perception in similar real and virtual environments. *ACM Transactions on Applied Perception (TAP)*, 2(1):35–45.
- Salmoni, A. W., Schmidt, R. A., and Walter, C. B. (1984). Knowledge of results and motor learning: a review and critical reappraisal. *Psychological bulletin*, 95(3):355.
- Sargent, R. G. (1998). Verification and validation of simulation models. In *Proceedings of the 30th conference on Winter simulation*, pages 121–130. IEEE Computer Society Press.
- Schmidt, R. and Lee, T. (2005). *Motor control and learning: A behavioral emphasis*. Human Kinetics, Champaign, IL, 4th ed edition.
- Schmidt, R. A. (1975). A schema theory of discrete motor skill learning. *Psychological review*, 82(4):225.

- Schuemie, M. J., Van Der Straaten, P., Krijn, M., and Van Der Mast, C. A. (2001). Research on presence in virtual reality: A survey. *CyberPsychology & Behavior*, 4(2):183–201.
- Schwebel, D. C., Gaines, J., and Severson, J. (2008). Validation of virtual reality as a tool to understand and prevent child pedestrian injury. *Accident Analysis & Prevention*, 40(4):1394–1400.
- Sebrechts, M. M. (2000). Using virtual environments for training of spatial navigation. *Human Systems IAC Gateway*, 11(2):1–4.
- Seo, K. (2004). A mysterious flight trajectory of a rugby ball. *Japanese Society of Mechanical Engineering News*, 15(2):4–5.
- Seo, K., Kobayashi, O., and Murakami, M. (2006). Multi-optimisation of the screw kick in rugby by using a genetic algorithm. *Sports Engineering*, 9(2):87–96.
- Seo, K., Kobayashi, O., Murakami, M., Yorita, D., Nagai, H., and Asai, K. (2010). Simulation of the trajectory of a punted rugby ball taking into account the asymmetrical pressure distribution caused by the seams. *Journal of visualization*, 13(2):97–105.
- Seo, K., Kobayashi, O., Sakamoto, S., and Murakami, M. (2004). Behavior of a rugby ball during flight. *Nihon Kikai Gakkai Nenji Taikai Koen Ronbunshu*, 7:45–46.
- Seyama, J. and Nagayama, R. S. (2007). The Uncanny Valley: Effect of Realism on the Impression of Artificial Human Faces. *Presence: Teleoperators and Virtual Environments*, 16(4):337–351.
- Seymour, N. E., Gallagher, A. G., Roman, S. A., OBrien, M. K., Bansal, V. K., Andersen, D. K., and Satava, R. M. (2002). Virtual reality training improves operating room performance: results of a randomized, double-blinded study. *Annals of surgery*, 236(4):458.

Shaw, D. E., Deneroff, M. M., Dror, R. O., Kuskin, J. S., Larson, R. H., Salmon, J. K., Young, C., Batson, B., Bowers, K. J., Chao, J. C., Eastwood, M. P., Gagliardo, J., Grossman, J. P., Ho, C. R., Ierardi, D. J., Kolossváry, I., Klepeis, J. L., Layman, T., McLeavey, C., Moraes, M. A., Mueller, R., Priest, E. C., Shan, Y., Spengler, J., Theobald, M., Towles, B., and Wang, S. C. (2008). Anton, a special-purpose machine for molecular dynamics simulation. *Commun. ACM*, 51(7):91–97.

Shibata, T., Kim, J., Hoffman, D. M., and Banks, M. S. (2011). Visual discomfort with stereo displays: effects of viewing distance and direction of vergence-accommodation conflict. In *Proceedings of SPIE*, volume 7863, pages 1–9. NIH Public Access.

sidvind.com (2014). Skybox Tutorial. http://sidvind.com/wiki/Skybox_tutorial, Date Accessed: 02/01/14.

Siemon, A., Wegener, R., Bader, F., Hieber, T., and Schmid, U. (2009). Video Games can Improve Performance in Sports. An Empirical Study with Wii Sports Bowling. In Althoff, K.-D., Back, K., and Reichle, M., editors, *KI 2009: Proceedings of the 32nd Annual Conference on Artificial Intelligence*, number September, page 74, Paderborn, Germany.

Simon, H. A. and Chase, W. G. (1973). Skill in chess: Experiments with chess-playing tasks and computer simulation of skilled performance throw light on some human perceptual and memory processes. *American scientist*, 61(4):394–403.

Skill-Tec (2014). Skill-Tec Rugby Ball. <http://store.r80rugby.com/Skill-Tec\%20Skills\%20Training\%20System.pdf>, Date Accessed: 02/01/14.

Slater, M. and Garau, M. (2007). The use of questionnaire data in presence studies: do not seriously likert. *Presence: Teleoperators and Virtual Environments*, 16(4):447–456.

- Slater, M., Steed, A., and Chrysanthou, Y. (2002). Computer graphics and virtual environments-from realism to real-time.
- Slater, M., Usoh, M., and Steed, A. (1994). Depth of presence in virtual environments. *Presence*, 3(2):130–144.
- Spathopoulos, V. M. (2011). A spreadsheet model for soccer ball flight mechanics simulation. *Computer Applications in Engineering Education*, 19(3):508–513.
- Squad (2014). Kerbal Space Program. <https://kerbalspaceprogram.com/about.php>, Date Accessed: 06/12/13.
- Starkes, J. (2000). The road to expertise: Is practice the only determinant? *International Journal of Sport Psychology*, 31(4):431–451.
- Starkes, J. L., Deakin, J. M., Allard, F., Hodges, N. J., and Hayes, A. (1996). Deliberate practice in sports: What is it anyway? In *The road to excellence: The acquisition of expert performance in the arts and sciences, sports, and games*, pages 81–106.
- Steinicke, F., Bruder, G., Hinrichs, K., Lappe, M., Ries, B., and Interrante, V. (2009). Transitional environments enhance distance perception in immersive virtual reality systems. In *Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization*, pages 19–26. ACM.
- Stereo Optical Company, Inc. (2014). Randot SO-002. Date Accessed: 02/01/14.
- Syamsuddin, M. R. and Kwon, Y. M. (2010). Research on virtual world and real world integration for batting practice. In *2010 International Symposium on Ubiquitous Virtual Reality*, pages 44–47. IEEE.
- Syamsuddin, M. R. and Kwon, Y. M. (2011). Simulation of baseball pitching and hitting on virtual world. In *2011 International Conference on Complex, Intelligent and Software Intensive Systems (CISIS)*, pages 663–667. IEEE.

- Takahashi, N. and Kobayashi, O. (2006). Auto-rotation on short axis of rugby ball. *Nihon Kikai Gakkai Nenji Taikai Koen Ronbunshu*, 2:123–124.
- The Hartree Centre (2014). The Hartree Centre at the Science and Technology Facilities Council, Daresbury, UK. <http://www.stfc.ac.uk/hartree/default.aspx>, Date Accessed: 03/01/14.
- Thomson, J. A. (1983). Is continuous visual monitoring necessary in visually guided locomotion? *Journal of Experimental Psychology: Human Perception and Performance*, 9(3):427.
- Tremblay, L. and Proteau, L. (1998). Specificity of practice: The case of powerlifting. *Research quarterly for exercise and sport*, 69(3):284–289.
- Triesch, J., Ballard, D. H., Hayhoe, M. M., and Sullivan, B. T. (2003). What you see is what you need. *Journal of vision*, 3(1).
- Ubisoft (2014). From Dust. <http://from-dust.ubi.com/from-dust/en-GB/home/index.aspx>, Date Accessed: 06/12/13.
- Ukai, K. (2007). Visual fatigue caused by viewing stereoscopic images and mechanism of accommodation. In *Proceedings of the First International Symposium on University Communication*, volume 1, pages 176–179.
- Ukai, K. and Howarth, P. A. (2008). Visual fatigue caused by viewing stereoscopic motion images: Background, theories, and observations. *Displays*, 29(2):106–116.
- Van Erp, J. B., Saturday, I., and Jansen, C. (2006). Application of tactile displays in sports: where to, how and when to move. In *Proceedings of the Eurohaptics International conference*.
- Vance, A., Buick, J., and Livesey, J. (2012). Aerodynamics of a rugby ball. *Transactions of the ASME-E-Journal of Applied Mechanics*, 79(2):021020.

- Vickers, J. N. (2007). Perception, cognition and decision training: the quiet eye in action.
- Vickers, J. N. and Adolphe, R. M. (1997). Gaze behaviour during a ball tracking and aiming skill. *International Journal of Sports Vision*, 4(1):18–27.
- Vignais, N., Bideau, B., Craig, C., Brault, S., Multon, F., Delamarche, P., and Kulpa, R. (2009a). Does the level of graphical detail of a virtual handball thrower influence a goalkeepers motor response. *Journal of Sports Science and Medicine*, 8:501–508.
- Vignais, N., Bideau, B., Kulpa, R., Craig, C., Brault, S., and Multon, F. (2009b). Virtual environments for sport analysis: Perception-action coupling in handball goal-keeping. *International Journal of Virtual Reality*, 8(4):43–8.
- Vignais, N., Kulpa, R., Craig, C., and Bideau, B. (2010a). Virtual thrower versus real goalkeeper: the influence of different visual conditions on performance. *Presence: Teleoperators and Virtual Environments*, 19(4):281–290.
- Vignais, N., Kulpa, R., Craig, C., Brault, S., Multon, F., and Bideau, B. (2010b). Influence of the graphical levels of detail of a virtual thrower on the perception of the movement. *Presence: Teleoperators and Virtual Environments*, 19(3):243–252.
- Vishwanath, D., Girshick, A. R., and Banks, M. S. (2005). Why pictures look right when viewed from the wrong place. *Nature neuroscience*, 8(10):1401–1410.
- Visionair (2014). <http://www.infra-visionair.eu/>, Date Accessed: 03/01/14.
- VR Media S.r.l. (2014). XVR Documentation: CVmVRPN Class Reference. http://www.vrmedia.it/Docs/class_c_vm_v_r_p_n.html, Date Accessed: 02/01/14.
- Wann, J. P. and Mon-Williams, M. (1997). Health issues with virtual reality displays: what we do know and what we don't. *ACM SIGGRAPH Computer Graphics*, 31(2):53–57.

- Watt, S. J., Akeley, K., Ernst, M. O., and Banks, M. S. (2005a). Focus cues affect perceived depth. *Journal of Vision*, 5(10).
- Watt, S. J., Akeley, K., Girshick, A. R., and Banks, M. S. (2005b). Achieving near-correct focus cues in a 3d display using multiple image planes. In *Proc. SPIE*, volume 5666, pages 393–401.
- Waxberg, S., Goodell, K., Avgerinos, D., Schwaitzberg, S., and Cao, C. (2004). Evaluation of physical versus virtual surgical training simulators. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 48, pages 1675–1679. SAGE Publications.
- Weisstein, E. W. (2014a). Linear equation. <http://mathworld.wolfram.com/LinearEquation.html>, Date Accessed: 02/01/14.
- Weisstein, E. W. (2014b). Newton's laws. <http://scienceworld.wolfram.com/physics/NewtonsLaws.html>, Date Accessed: 02/01/14.
- Weisstein, E. W. (2014c). Runge-kutta method. <http://mathworld.wolfram.com/Runge-KuttaMethod.html>, Date Accessed: 02/01/14.
- Welsh Assembly Government (2009). Welsh Health Survey 2009. <http://wales.gov.uk/docs/statistics/2010/100915healthsurvey09en.pdf>, Date Accessed: 02/01/14.
- Willemsen, P., Colton, M. B., Creem-Regehr, S. H., and Thompson, W. B. (2004). The effects of head-mounted display mechanics on distance judgments in virtual environments. In *Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, pages 35–38. ACM.
- Willemsen, P., Colton, M. B., Creem-Regehr, S. H., and Thompson, W. B. (2009). The effects of head-mounted display mechanical properties and field of view on distance

- judgments in virtual environments. *ACM Transactions on Applied Perception (TAP)*, 6(2):8.
- Willemsen, P., Gooch, A. A., Thompson, W. B., and Creem-Regehr, S. H. (2008). Effects of stereo viewing conditions on distance perception in virtual environments. *Presence: Teleoperators and Virtual Environments*, 17(1):91–101.
- Williams, B., Johnson, D., Shores, L., and Narasimham, G. (2008). Distance perception in virtual environments. In *Proceedings of the 5th symposium on Applied perception in graphics and visualization*, pages 193–193. ACM.
- Witmer, B. G., Jerome, C. J., and Singer, M. J. (2005). The factor structure of the presence questionnaire. *Presence: Teleoperators and Virtual Environments*, 14(3):298–312.
- Witmer, B. G. and Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and virtual environments*, 7(3):225–240.
- Wloka, M. M. (1995). Lag in multiprocessor virtual reality. *Presence: Teleoperators and Virtual Environments*, 4(1):50–63.
- Woodworth, R. S. (1899). Accuracy of voluntary movement. *The Psychological Review: Monograph Supplements*, 3(3).
- Xu, S., Song, P., Chin, C. L., Chua, G. G., Huang, Z., and Rahardja, S. (2009). Tennis space: an interactive and immersive environment for tennis simulation. In *Image and Graphics, 2009. ICIG'09. Fifth International Conference on*, pages 652–657. IEEE.
- Yano, S., Emoto, M., and Mitsuhashi, T. (2004). Two factors in visual fatigue caused by stereoscopic hdtv images. *Displays*, 25(4):141–150.

Zaal, F. T. and Bootsma, R. J. (2011). Virtual reality as a tool for the study of perception-action: the case of running to catch fly balls. *Presence: Teleoperators and Virtual Environments*, 20(1):93–103.

Zaal, F. T. and Michaels, C. F. (2003). The information for catching fly balls: Judging and intercepting virtual balls in a CAVE. *Journal of Experimental Psychology: Human Perception and Performance*, 29(3):537.

List of Terms

Accommodation Accommodation is the process of the eye muscles changing the focal length of the lens in the eye in order to focus on an object.

Anaglyphic The use of differently coloured lenses for each eye in a pair of glasses to restrict the portions of an image entering the left and right eyes.

ANOVA Analysis of Variance, a method of statistical analysis that compares the mean values of data sets.

API Application Programming Interface.

Arcminutes/Arcseconds A unit of angular measurement in optometry: an arcminute is 1/60th of a degree, and an arcsecond is 1/60th of an arcminute (or 1/3600th of a degree).

ART Advanced Realtime Tracking GmbH, Weilheim, Germany. A company producing a range of motion capture solutions. <http://www.ar-tracking.com/home/>.

Azimuth In a spherical coordinate system around a person, the azimuth encircles the person horizontally, parallel to the ground.

Binocular Cues Depth cues in the human visual system that require both of the eyes to interpret correctly: *binocular disparity* and *vergence*.

Binocular Disparity The physical separation between a person's eyes results in each eye viewing a different image, the difference in the two images is the binocular disparity.

Cave Automatic Virtual Environment (CAVE) A CAVE is a group of large projected screens formed into a cube which the user stands within to be surrounded by the walls. CAVEs may include floor and ceiling screens, but are not always comprised of four walls (three are commonly used as the user faces away from the open wall).

Centre de Réalité Virtuelle de la Méditerranée (CRVM) Centre for the Virtual Reality of the Mediterranean at Université de la Méditerranée: Aix-Marseille II in Marseille, France. <http://www.realite-virtuelle.univmed.fr/en>.

CRP Corneal Reflection Pupillometer. A corneal reflection pupillometer is a device to digitally measure interpupillary distance: the device is placed in front of the subject's eyes and the vertically projected lines must be aligned to the centre of the subject's pupils to obtain a reading.

Cybersickness A form of motion sickness where the visual movement experienced by the user does not match the movement sensed by the vestibular system.

DLP Digital Light Processing.

Egocentric Distance A person's measurement of distance between an object and their own body.

FOV Field of View.

Ghosting Ghosting is when an on-screen image is partially reproduced twice with a slight offset, giving the appearance of a 'ghost' next to the main subject. It

can occur from crosstalk in stereoscopic systems, where part of the image leaks through the filter to the incorrect eye.

GIMP GNU Image Manipulation Program.

Haptic Device Hardware to simulate tactile sensations.

Head-Mounted Display (HMD) A head-mounted display is a device worn on the head that supports either a single screen or a pair of screens in front of the user's eyes.

Inattentive Blindness When a person is focused on a specific task, they may fail to notice an unexpected stimulus occurring in their field of vision.

Infitec Interference filter technology by Infitec GmbH, Ulm, Germany.

International Rugby Board (IRB) The International Rugby Board are responsible for maintaining the laws relating to the game of Rugby Union. The Laws of the Game are available online at <http://www.irblaws.com>.

Interpupillary Distance (IPD) Interpupillary distance is the measurement of the physical distance between a person's pupils; the distance is measured from the centre of each pupil as the person is positioned with their eyes facing directly forwards.

Latency Time delays experienced due to the passing of data between different hardware and software.

Likert Scale A psychometric scale for questionnaires; the responses to several items are grouped together to provide insight into the subject's attitudes.

Magnus Effect When a ball is flying through the air, the direction of spin on the ball will cause it to curve from its flight path.

MiddleVR Middleware for connecting hardware and software by I'm in VR.

ML Missed Left.

Monocular Cues Depth cues in the human visual system which do not require both of the eyes to interpret correctly: *accommodation*, *motion parallax*, *linear perspective* (where parallel lines appear to become closer in the distance), *shading*, *height in the visual field*, *occlusion* (closer objects are visibly in front of further ones), *blur* (objects further away are more blurred).

Motion Parallax When focussing on an object and moving the head from side to side, objects that are closer appear to move less than objects that are further away.

MR Missed Right.

OT On Target.

Presence Feelings of “being there” within a VE; a psychological concept in which the user of a VE feels as though the environment exists in reality and they are truly present within it.

PST Personal Space Tracker. Small tracking device by Personal Space Technologies. More information is available at <http://ps-tech.com/tracking/pst-55/>.

Random Dot Stereogram A random dot stereogram is an image comprised of random dots that when viewed with stereoscopic glasses produces the illusion of depth; they can be used to assess a subject’s stereoacuity by varying the amount of binocular disparity across a series of images.

Retroreflective Material that is capable of reflecting a light back to its source from any angle.

Runge-Kutta 4th Order (RK4) Runge-Kutta are a group of methods for approximating the solution to ordinary differential equations using iterative steps. The 4th

order method contains four incremental steps used to approximate the next value from a given initial value.

Stereoacuity Stereoacuity (or stereoscopic acuity) is the measurement of the smallest amount of binocular disparity a subject is capable of viewing. Normal stereoacuity falls in the range of 20–500 arcseconds (0.3–8.3 arcminutes).

Stereoscopy The illusion of depth is created in a flat image by presenting each eye with an image from a slightly different angle (stereopsis).

UDP User Datagram Protocol. Internet Standard RFC 768 for User Datagram Protocol, available at <http://tools.ietf.org/html/rfc768>..

Unity 3D Game development software by Unity Technologies.

Vergence Vergence describes the movement of the two eyes simultaneously: *convergence* is rotating the two eye inwards to view close objects, and *divergence* is rotating the two eyes outwards until they are parallel, to view objects in the distance.

VERST Virtual Environment for Rugby Skills Training.

Vicon Vicon, Denver, USA. A company producing a range of optical tracking solutions. <http://www.vicon.com/>.

Virtools Virtual environment development software by Dassault Systemes..

Virtual Engineering Centre (VEC) A facility run by Liverpool University at the Science and Technology Facilities Council in Daresbury, UK.

Virtual Environment (VE) A virtual environment is a collection of hardware and software components, formed to immerse a user into a virtual interactive scene.

VO Virtual Opponent.

VR Virtual Reality.

VRPN The Virtual Reality Peripheral Network (open source), a server which is able to receive data streamed from many different tracking devices and process it into a single format, compatible with many software packages that can poll the server for data as required.

XVR eXtreme Virtual Reality, virtual environment development software by VR Media S.r.l..

Chapter 7

Appendix

Additional Materials:

7.1 Software Evaluation Tables

7.2 IRB Regulations: The Pitch

7.3 IRB Regulations: The Ball

7.4 Physics Engine Code (XVR)

7.4.1 Function: `Integrate()`

7.4.2 Function: `Rigid1()`

7.4.3 Function: `Rigid2()`

7.5 Experiment 1

7.5.1 Full Specifications

7.5.2 Participant Guidelines and Consent Form

7.6 Experiment 2

7.6.1 Full Specifications

7.6.2 Configuration Files for Verbal Task

7.6.3 Configuration Files for Throwing Task

7.6.4 Participant Guidelines and Consent Form

7.6.4.1 Throwing Task

7.6.4.2 Verbal Task

7.6.5 Presence Questionnaire

7.6.6 Presence Questionnaire Full Results

7.7 Poster Submitted to JVRC 2011

7.1 Software Evaluation Tables

Table 7.1: Software for interfacing between the hardware and the VE.

Name	Developer	Positive	Negative
Conduit	Mechdyne	<ul style="list-style-type: none"> • Custom support from Mechdyne 	<ul style="list-style-type: none"> • Not free • For distributing data to multiple displays or workstations
Trackd	Mechdyne	<ul style="list-style-type: none"> • 12 months professional maintenance 	<ul style="list-style-type: none"> • Not free
VR Juggler	Iowa State University's Virtual Reality Applications Center	<ul style="list-style-type: none"> • Free • Large range of compatible devices 	
VRPN	Department of Computer Science, University of North Carolina at Chapel Hill	<ul style="list-style-type: none"> • Free • Large range of compatible devices • Compatibility built into the PS Track software 	
MiddleVR	I'm in VR	<ul style="list-style-type: none"> • Large range of compatible devices • Easy to configure complex hardware • Compatible with VRPN 	<ul style="list-style-type: none"> • Not free to use full version

Table 7.2: Software for creating the virtual scene.

Name	Developer	Positive	Negative	Compatibility
3DVIA Virtools	Dassault Systemes	<ul style="list-style-type: none"> • Stereoscopy • Windows 7 	<ul style="list-style-type: none"> • Not free 	<ul style="list-style-type: none"> • OpenGL • VRPN
Avango	Fraunhofer IAIS	<ul style="list-style-type: none"> • Free • Windows 7 • Stereoscopy • Rapid development • Runs on top of Open Scene Graph 	<ul style="list-style-type: none"> • Doesn't have own scene graph module • Although a 3D application engine, doesn't have game engine functions built in • Developed mainly for Linux, and only ported to Windows 	<ul style="list-style-type: none"> • C++ • Python • VRPN
CAVELib	Mechdyne	<ul style="list-style-type: none"> • Stereoscopy • Windows 7 	<ul style="list-style-type: none"> • Not free • Unsure if a CAVE will be used 	<ul style="list-style-type: none"> • C • C++ • OpenGL • Trackd • VRPN
EON Suite	EON Reality	<ul style="list-style-type: none"> • Stereoscopy • Rapid development • Windows 7 	<ul style="list-style-type: none"> • Not free • Suited better to industrial simulation 	<ul style="list-style-type: none"> • Trackd
Maverik	AIG at Manchester University	<ul style="list-style-type: none"> • Free • Windows 7 	<ul style="list-style-type: none"> • Not known if there is support for stereoscopic content • No longer in development 	<ul style="list-style-type: none"> • C++ • VRPN
Ogre3D	Freeware	<ul style="list-style-type: none"> • Free • Windows 7 	<ul style="list-style-type: none"> • Not known if there is support for stereoscopic content 	<ul style="list-style-type: none"> • C++ • VRPN

Name	Developer	Positive	Negative	Compatibility
Open Scene Graph	Open Source	<ul style="list-style-type: none"> • Free • Windows 7 • Stereoscopic 		<ul style="list-style-type: none"> • C++ • OpenGL • VRPN
Panda 3D	Carnegie Mellon University	<ul style="list-style-type: none"> • Free • Windows 7 • Stereoscopic • Inbuilt physics 	<ul style="list-style-type: none"> • Aimed at game development rather than VEs 	<ul style="list-style-type: none"> • C++ • Python • VRPN
Syzygy	University of Illinois at Urbana-Champaign	<ul style="list-style-type: none"> • Free 	<ul style="list-style-type: none"> • Not known if compatible with Windows 7 • Not known if there is support for stereoscopic content 	<ul style="list-style-type: none"> • C++ • Python • Trackd • VRPN
Unity 3D	Unity Technologies	<ul style="list-style-type: none"> • Free and Pro versions available • Windows 7 • Inbuilt physics engine 	<ul style="list-style-type: none"> • Stereoscopic content available with Pro version only • Aimed at game development rather than VEs 	Scripting Languages: <ul style="list-style-type: none"> • C# • Boo • Unity
Vizard Virtual Reality Toolkit	WorldViz	<ul style="list-style-type: none"> • Windows 7 • Stereoscopic 	<ul style="list-style-type: none"> • Not free to use; \$3,990 (approx 2500 for academic license) 	<ul style="list-style-type: none"> • C++ • Python
XVR	VR Media S.r.l.	<ul style="list-style-type: none"> • Free • Windows 7 • Stereoscopic • Developed for virtual reality • Can run platform-independently 		<ul style="list-style-type: none"> • OpenGL • VRPN

7.2 IRB Regulations: The Pitch




Figure 7.1: Ground specifications according to law 1 of the IRB game laws

According to the IRB Laws of the Game, available online at <http://www.irblaws.com>, the pitch (see figure 7.1) used for the game must be:

Law 1.2 Required Dimensions for the Playing Enclosure

- (a) The field of play does not exceed 100 metres in length and 70 metres in width. Each in-goal does not exceed 22 metres in length and 70 metres in width.
- (b) The length and breadth of the playing area are to be as near as possible to the dimensions indicated. All the areas are rectangular.
- (c) The distance from the goal line to the dead ball line should be not less than 10 metres where practicable.

Law 1.4 Dimensions for Goal Post and Crossbar

- (a) The distance between the two goal posts is 5.6 metres.
- (b) The crossbar is placed between the two goal posts so that its top edge is 3.0 metres from the ground.
- (c) The minimum height of the goal posts is 3.4 metres.
- (d) When padding is attached to the goal posts the distance from the goal line to the external edge of the padding must not exceed 300 mm.

7.3 IRB Regulations: The Ball




Figure 7.2: Ball specifications according to law 2 of the IRB game laws.

According to the IRB Laws of the Game, available online at <http://www.irblaws.com>, the ball (see figure 7.2) used for the game must be:

Law 2.1 Shape

The ball must be oval and made of four panels.

Law 2.2 Dimensions

Length in line	280 - 300 mm
Circumference (end to end)	740 - 770 mm
Circumference (in width)	580 - 620 mm

Law 2.4 Weight

410 - 460 grams

7.4 Physics Engine Code (XVR)

7.4.1 Function: Integrate()

```
1 function Integrate(state, dt)
2 {
3     k1 = Rigid1(state);
4     k2 = Rigid2(state, dt*0.50, k1);
5     k3 = Rigid2(state, dt*0.50, k2);
6     k4 = Rigid2(state, dt, k3);
7
8     // x
9     var d1dt = 1.00/6.00 * (k1[0] + (2*k2[0]) + (2*k3[0]) + k4[0]);
10    var d2dt = 1.00/6.00 * (k1[1] + (2*k2[1]) + (2*k3[1]) + k4[1]);
11    // y
12    var d3dt = 1.00/6.00 * (k1[2] + (2*k2[2]) + (2*k3[2]) + k4[2]);
13    var d4dt = 1.00/6.00 * (k1[3] + (2*k2[3]) + (2*k3[3]) + k4[3]);
14    // z
15    var d5dt = 1.00/6.00 * (k1[4] + (2*k2[4]) + (2*k3[4]) + k4[4]);
16    var d6dt = 1.00/6.00 * (k1[5] + (2*k2[5]) + (2*k3[5]) + k4[5]);
17
18    // pos = pos + vel*time
19    // vel = vel + acc*time
20    ball[0] = state[0] + d1dt*dt; // x pos
21    ball[1] = state[1] + d2dt*dt; // x vel
22    ball[2] = state[2] + d3dt*dt; // y pos
23    ball[3] = state[3] + d4dt*dt; // y vel
24    ball[4] = state[4] + d5dt*dt; // z pos
25    ball[5] = state[5] + d6dt*dt; // z vel
26
27    return ball;
28 }
```

Figure 7.3: Function to approximate the next position of the ball using RK4.

7.4.2 Function: Rigid1 ()

```
1 function Rigid1(initial_ball)
2 {
3     // x,y,z velocities
4     temp.x = initial_ball[1];
5     temp.y = initial_ball[3];
6     temp.z = initial_ball[5];
7
8     // calculate the current velocity magnitude
9     v = sqrt(temp.x^2 + temp.y^2 + temp.z^2);
10
11    // formulae for rugby ball (y and z swapped to fit with XVR
12    // coord system)
13    // x
14    output_dball1[0] = temp.x;
15    output_dball1[1] = (-v)*((Kd * temp.x) - (Kl * ((cos(gamma) *
16    temp.z) + (sin(gamma) * temp.y))));
17    // y
18    output_dball1[2] = temp.y;
19    output_dball1[3] = (-g)-v*((Kd * temp.z) + (Kl * cos(gamma) *
20    temp.x));
21    // z
22    output_dball1[4] = temp.z;
23    output_dball1[5] = (-v)*((Kd * temp.y) - (Kl * sin(gamma) * temp
24    .x));
25
26    return output_dball1;
27 }
```

Figure 7.4: Function to evaluate k_1 using equations 3.22, 3.23 and 3.24.

7.4.3 Function: Rigid2 ()

```
1 function Rigid2(initial_ball, dt, e)
2 {
3     // x
4     new_ball[0] = initial_ball[0] + e[0]*dt;
5     new_ball[1] = initial_ball[1] + e[1]*dt;
6     // y
7     new_ball[2] = initial_ball[2] + e[2]*dt;
8     new_ball[3] = initial_ball[3] + e[3]*dt;
9     // z
10    new_ball[4] = initial_ball[4] + e[4]*dt;
11    new_ball[5] = initial_ball[5] + e[5]*dt;
12
13    // x,y,z velocities
14    temp.x = new_ball[1];
15    temp.y = new_ball[3];
16    temp.z = new_ball[5];
17
18    // calculate the current velocity magnitude
19    v = sqrt(temp.x^2 + temp.y^2 + temp.z^2);
20
21    // formulae for rugby ball (y and z swapped to fit with XVR
22    // coord system)
23    // x
24    output_dball2[0] = temp.x;
25    output_dball2[1] = (-v) * ((Kd * temp.x) - (Kl * ((cos(gamma) *
26    temp.z) + (sin(gamma) * temp.y))));
27    // y
28    output_dball2[2] = temp.y;
29    output_dball2[3] = (-g)-(v) * ((Kd * temp.z) + (Kl * cos(gamma)
30    * temp.x));
31    // z
32    output_dball2[4] = temp.z;
33    output_dball2[5] = (-v) * ((Kd * temp.y) - (Kl * sin(gamma) *
34    temp.x));
35
36    return output_dball2;
37 }
```

Figure 7.5: Function to evaluate k_2 , k_3 and k_4 using equations 3.22, 3.23 and 3.24.

7.5 Experiment 1

7.5.1 Full Specification

Table 7.3: Full Configuration Procedure of VERST for Experiment 1

Subject	Trial No.	Subject Position	Stereo/No Stereo	Floor/No Floor	Target Distance	Repetitions at Different Angles	Total No. Trials
1	001 - 035	A	S	F	× 7	× 5	35
	036 - 070	A	S	NF	× 7	× 5	70
	071 - 105	A	NS	F	× 7	× 5	105
	106 - 140	A	NS	NF	× 7	× 5	140
	141 - 175	B	S	F	× 7	× 5	175
	176 - 210	B	S	NF	× 7	× 5	210
	211 - 245	B	NS	F	× 7	× 5	245
	246 - 280	B	NS	NF	× 7	× 5	280
2	001 - 035	B	S	F	× 7	× 5	315
	036 - 070	B	S	NF	× 7	× 5	350
	071 - 105	B	NS	F	× 7	× 5	385
	106 - 140	B	NS	NF	× 7	× 5	420
	141 - 175	A	S	F	× 7	× 5	455
	176 - 210	A	S	NF	× 7	× 5	490
	211 - 245	A	NS	F	× 7	× 5	525
	246 - 280	A	NS	NF	× 7	× 5	560
3	001 - 035	A	NS	NF	× 7	× 5	595
	036 - 070	A	NS	F	× 7	× 5	630
	071 - 105	A	S	NF	× 7	× 5	665
	106 - 140	A	S	F	× 7	× 5	700
	141 - 175	B	NS	NF	× 7	× 5	735

Subject	Trial No.	Subject Position	Stereo/No Stereo	Floor/No Floor	Target Distance	Repetitions at Different Angles	Total No. Trials
	176 - 210	B	NS	F	× 7	× 5	770
	211 - 245	B	S	NF	× 7	× 5	805
	246 - 280	B	S	F	× 7	× 5	840
4	001 - 035	B	NS	NF	× 7	× 5	875
	036 - 070	B	NS	F	× 7	× 5	910
	071 - 105	B	S	NF	× 7	× 5	945
	106 - 140	B	S	F	× 7	× 5	980
	141 - 175	A	NS	NF	× 7	× 5	1015
	176 - 210	A	NS	F	× 7	× 5	1050
	211 - 245	A	S	NF	× 7	× 5	1085
	246 - 280	A	S	F	× 7	× 5	1120

7.5.2 Participant Guidelines and Consent Form

Participant Guidelines

An Experiment to Investigate Distance Perception in a Virtual Environment

Researcher: Helen Miles, Bangor University, North Wales, UK.

Purpose of the Study

The purpose of this experiment is to determine whether objects in a virtual scene shown in the CAVE environment at CRVM appear to be at the correct distance. We are also interested in the effect of using a stereoscopic display and a floor screen on the perception of distance.

Procedure

You will be asked to throw a real rugby ball towards a virtual target in a number of different conditions: with or without stereoscopic display, with or without a floor screen and at different distances away from the screen. The target will be an archery-style bull's-eye target, and you will have to throw for targets at several different distances. The real ball that you will be throwing is attached to the ceiling with a bungee cord to prevent damage to the equipment.

Before the experiment begins, you will be allowed some time to warm up; you will be observed aiming at a real-world target. You will also be allowed a few practice throws with the tethered ball. There will be a short visual test before the experiment begins; then the experiment will take approximately 1 hour to complete. If you consent, you may be filmed for parts of this experiment for analysis later on.

During the experiment you will be shown the virtual target, you will then have 5 seconds to make an attempt to hit the target with the ball. A 'beep' will signify the start of the trial, from that time you will have 5 seconds to make an attempt to hit the target. If you failed to make an attempt in the 5 seconds or the tracking system fails to capture the ball, a buzzer will sound and you will be allowed one further attempt at the trial.

Although we have a mathematical model to predict the flight of the rugby ball, you will not see the ball and will not be told whether you have hit the target.

Confidentiality

Your data will remain confidential and will be used only for the research purposes of this experiment.

Participation and Withdrawal

Your participation in this experiment is voluntary and you may choose to withdraw at any time without giving a reason.

Participant Consent Form

An Experiment to Investigate Distance Perception in a Virtual Environment
Researcher: Helen Miles, Bangor University, North Wales, UK.

Name:	Age:	Gender:	M	F
Please give some details of your experience playing rugby (e.g. years playing, playing position):				
Do you wear glasses/corrective lenses?				
Do you have any history of eyesight problems?				
Are you colour-blind?				
Are you left or right-handed?				
If yes to any of the above, please give details:				
Do you consent to being filmed during this experiment?				

I have read the participant guidelines sheet for the above study and had the opportunity to ask questions. I understand that my voluntary participation means I may withdraw at any time without giving any reason. I agree to participate in this study.

Participant Signature

Date.....

7.6 Experiment 2

7.6.1 Full Specification

Table 7.4: Full Configuration Procedure of VERST for Experiment 2

Subject	Trial No.	Task (Verbal/Throwing)	Screen Type	Subject Position	Stereo/No Stereo	Target Distance	Repetitions at Different Angles	Total No. Trials	
1	001 - 018	V	F	N	S	× 6	× 3	18	
	019 - 036	V	F	M	S	× 6	× 3	36	
	037 - 054	V	F	F	S	× 6	× 3	54	
	055 - 072	V	F	N	NS	× 6	× 3	72	
	073 - 090	V	F	M	NS	× 6	× 3	90	
	091 - 108	V	F	F	NS	× 6	× 3	108	
	109 - 126	T	F	F	S	× 6	× 3	126	
	127 - 144	T	F	M	S	× 6	× 3	144	
	145 - 162	T	F	M	NS	× 6	× 3	162	
	163 - 180	T	F	F	NS	× 6	× 3	180	
	181 - 198	V	C	N	NS	× 6	× 3	198	
	199 - 216	V	C	M	NS	× 6	× 3	216	
	217 - 234	V	C	F	NS	× 6	× 3	234	
	235 - 252	V	C	N	S	× 6	× 3	252	
	253 - 270	V	C	M	S	× 6	× 3	270	
	271 - 288	V	C	F	S	× 6	× 3	288	
	2	001 - 018	V	F	M	S	× 6	× 3	306
		019 - 036	V	F	F	S	× 6	× 3	324
037 - 054		V	F	N	NS	× 6	× 3	342	
055 - 072		V	F	M	NS	× 6	× 3	360	
073 - 090		V	F	F	NS	× 6	× 3	378	

Subject	Trial No.	Task (Verbal/Throwing)	Screen Type	Subject Position	Stereo/No Stereo	Target Distance	Repetitions at Different Angles	Total No. Trials
	091 - 108	V	F	N	S	× 6	× 3	396
	109 - 126	T	F	M	S	× 6	× 3	414
	127 - 144	T	F	F	S	× 6	× 3	432
	145 - 162	T	F	M	NS	× 6	× 3	450
	163 - 180	T	F	F	NS	× 6	× 3	468
	181 - 198	V	C	F	S	× 6	× 3	486
	199 - 216	V	C	N	NS	× 6	× 3	504
	217 - 234	V	C	M	NS	× 6	× 3	522
	235 - 252	V	C	F	NS	× 6	× 3	540
	253 - 270	V	C	N	S	× 6	× 3	558
	271 - 288	V	C	M	S	× 6	× 3	576
3	001 - 018	V	C	M	S	× 6	× 3	594
	019 - 036	V	C	F	S	× 6	× 3	612
	037 - 054	V	C	N	NS	× 6	× 3	630
	055 - 072	V	C	M	NS	× 6	× 3	648
	073 - 090	V	C	F	NS	× 6	× 3	666
	091 - 108	V	C	N	S	× 6	× 3	684
	109 - 126	V	F	F	S	× 6	× 3	702
	127 - 144	V	F	N	NS	× 6	× 3	720
	145 - 162	V	F	M	NS	× 6	× 3	738
	163 - 180	V	F	F	NS	× 6	× 3	756
	181 - 198	V	F	N	S	× 6	× 3	774
	199 - 216	V	F	M	S	× 6	× 3	792
	217 - 234	T	F	F	S	× 6	× 3	810
	235 - 252	T	F	M	NS	× 6	× 3	828
	253 - 270	T	F	M	S	× 6	× 3	846

Subject	Trial No.	Task (Verbal/Throwing)	Screen Type	Subject Position	Stereo/No Stereo	Target Distance	Repetitions at Different Angles	Total No. Trials
	271 - 288	T	F	F	NS	× 6	× 3	864
4	001 - 018	V	F	N	NS	× 6	× 3	882
	019 - 036	V	F	M	NS	× 6	× 3	900
	037 - 054	V	F	F	NS	× 6	× 3	918
	055 - 072	V	F	N	S	× 6	× 3	936
	073 - 090	V	F	M	S	× 6	× 3	954
	091 - 108	V	F	F	S	× 6	× 3	972
	109 - 126	T	F	M	NS	× 6	× 3	990
	127 - 144	T	F	F	S	× 6	× 3	1008
	145 - 162	T	F	F	NS	× 6	× 3	1026
	163 - 180	T	F	M	S	× 6	× 3	1044
	181 - 198	V	C	N	S	× 6	× 3	1062
	199 - 216	V	C	M	S	× 6	× 3	1080
	217 - 234	V	C	F	S	× 6	× 3	1098
	235 - 252	V	C	N	NS	× 6	× 3	1116
	253 - 270	V	C	M	NS	× 6	× 3	1134
	271 - 288	V	C	F	NS	× 6	× 3	1152
5	001 - 018	V	C	F	NS	× 6	× 3	1170
	019 - 036	V	C	N	S	× 6	× 3	1188
	037 - 054	V	C	M	S	× 6	× 3	1206
	055 - 072	V	C	F	S	× 6	× 3	1224
	073 - 090	V	C	N	NS	× 6	× 3	1242
	091 - 108	V	C	M	NS	× 6	× 3	1260
	109 - 126	V	F	M	NS	× 6	× 3	1278
	127 - 144	V	F	F	NS	× 6	× 3	1296
	145 - 162	V	F	N	S	× 6	× 3	1314

Subject	Trial No.	Task (Verbal/Throwing)	Screen Type	Subject Position	Stereo/No Stereo	Target Distance	Repetitions at Different Angles	Total No. Trials
	163 - 180	V	F	M	S	× 6	× 3	1332
	181 - 198	V	F	F	S	× 6	× 3	1350
	199 - 216	V	F	N	NS	× 6	× 3	1368
	217 - 234	T	F	F	NS	× 6	× 3	1386
	235 - 252	T	F	M	S	× 6	× 3	1404
	253 - 270	T	F	F	S	× 6	× 3	1422
	271 - 288	T	F	M	NS	× 6	× 3	1440
6	001 - 018	V	C	M	NS	× 6	× 3	1458
	019 - 036	V	C	F	NS	× 6	× 3	1476
	037 - 054	V	C	N	S	× 6	× 3	1494
	055 - 072	V	C	M	S	× 6	× 3	1512
	073 - 090	V	C	F	S	× 6	× 3	1530
	091 - 108	V	C	N	NS	× 6	× 3	1548
	109 - 126	V	F	F	NS	× 6	× 3	1566
	127 - 144	V	F	N	S	× 6	× 3	1584
	145 - 162	V	F	M	S	× 6	× 3	1602
	163 - 180	V	F	F	S	× 6	× 3	1620
	181 - 198	V	F	N	NS	× 6	× 3	1638
	199 - 216	V	F	M	NS	× 6	× 3	1656
	217 - 234	T	F	M	S	× 6	× 3	1674
	235 - 252	T	F	F	S	× 6	× 3	1692
	253 - 270	T	F	F	NS	× 6	× 3	1710
	271 - 288	T	F	M	NS	× 6	× 3	1728

7.6.2 Configuration Files for Verbal Task

Figure 7.6 and figure 7.7 are examples of the two files used to prepare the verbal experiment. Figure 7.6 contains only information about the timings of the experiment, while figure 7.7 contains the participant ID number and the order they would complete the configurations. The files were read by line numbers, and by keeping each number on a separate line, it was not necessary to employ complex any token-splitting techniques.

```
1 wait between trials (s):  
2 1  
3  
4 time to complete trial (s):  
5 5  
6  
7 countdown between configuration changes (s):  
8 5  
9  
10 countdown before experiment starts (s):  
11 5
```

Figure 7.6: Trial timings are read from this file.

```
1 participant ID:  
2 1  
3  
4 configuration order:  
5 1  
6 2  
7 3  
8 4  
9 5  
10 6
```

Figure 7.7: Participant number is used to number the output files, while the configuration order is read into the program and interpreted to set order the configurations are called.

7.6.3 Configuration Files for Throwing Task

Figure 7.8 and figure 7.9 are examples of the files used to set up the throwing experiment. In addition to all of the settings contained in the verbal experiment, figure 7.8 contains settings regarding the tracker, and whether the virtual ball should be visible when thrown.

```
1 wait between trials (s):
2 1
3
4 time to complete trial (s):
5 5
6
7 countdown between configuration changes (s):
8 5
9
10 countdown before experiment starts (s):
11 5
12
13 tracker offset [xyz]:
14 0 0 0
15
16 capture window offset from participant's head position [xyz]:
17 -1
18
19 unity throwing axis (x:forward, y:up, z:left):
20 x
21
22 tracker height variable, amount to compare the height of the ball
   before it hits the virtual ground:
23 0
24
25 show ball (Y/N)?
26 N
```

Figure 7.8: Trial timings are read from this file, along with tracker settings.

```
1 participant ID:
2 1
3
4 configuration order:
5 2
6 1
7 3
8 4
```

Figure 7.9: Participant number is used to number the output files, while the configuration order is read into the program and interpreted to set order the configurations are called.

7.6.4 Participant Guidelines and Consent Form

7.6.4.1 Throwing Task

Participant Guidelines

An Experiment to Investigate Distance Perception in a Virtual Environment

Researcher: Helen Miles, Bangor University, North Wales, UK.

Purpose of the Study

The purpose of this experiment is to determine whether objects in a virtual scene shown in the large flat screen environment at the Hartree Centre appear to be at the correct distance. We are also interested in the effect of using a stereoscopic display and the user's distance from the screen on the perception of distance.

Procedure

You will be asked to throw a real rugby ball the distance of a virtual target in a number of different conditions: with or without stereoscopic display, throwing the ball in different directions and at different distances away from the screen. The target will be an archery-style bull's-eye target, and you will have to throw to targets at several different distances. The real ball that you will be throwing is attached to a wrist strap with a bungee cord to prevent damage to the equipment.

Before the experiment begins, you will be allowed some time to warm up; you will be observed aiming at a real-world target. You will also be allowed a few practice throws with the tethered ball. There will be a short visual test before the experiment begins; then the experiment will take approximately 30 minutes to complete. If you consent, you may be filmed in parts of this experiment for later analysis.

During the experiment you will be shown the virtual target, you will then have 5 seconds to make an attempt to hit the target with the ball. If you failed to make an attempt in the 5 seconds or the tracking system fails to capture the ball, the trial will be nullified. Although we have a physics engine to model the flight of the rugby ball based on your throw, it will not be visible for this experiment.

Confidentiality

Your data will remain confidential and will be used only for the research purposes of this experiment.

Participation and Withdrawal

Your participation in this experiment is voluntary and you may choose to withdraw at any time without giving a reason.

7.6.4.2 Verbal Task

Participant Guidelines

An Experiment to Investigate Distance Perception in a Virtual Environment

Researcher: Helen Miles, Bangor University, North Wales, UK.

Purpose of the Study

The purpose of this experiment is to determine whether objects in a virtual scene shown in the large flat and curved screen environments at the Hartree Centre appear to be at the correct distance. We are also interested in the effect of using a stereoscopic display and the user's distance from the screen on the perception of distance.

Procedure

You will be asked to dictate the distance of a virtual target in a number of different conditions: with or without stereoscopic display, and at different distances away from the screen. The target will be an archery-style bull's-eye target, and there will be targets at several different distances. This experiment will be performed twice, once in front of a flat screen and once in front of a curved screen.

There will be a short visual test before the experiment begins; then the experiment will take approximately 10–15 minutes to complete in each room (30 minutes total). If you consent, you may be photographed/filmed in parts of this experiment for later analysis.

During the experiment you will be shown the virtual target, you will then have 5 seconds to dictate the distance. If you failed to make an attempt in the 5 seconds, a null answer will be recorded.

Confidentiality

Your data will remain confidential and will be used only for the research purposes of this experiment.

Participation and Withdrawal

Your participation in this experiment is voluntary and you may choose to withdraw at any time without giving a reason.

Participant Consent Form

An Experiment to Investigate Distance Perception in a Virtual Environment

Researcher: Helen Miles, Bangor University, North Wales, UK.

Name:	Age:	Gender:	M	F		
Please give some details of your experience playing rugby (e.g. years playing, playing position):						
Do you wear glasses/corrective lenses?	Yes	No				
Do you have any history of eyesight problems?	Yes	No				
Are you colour-blind?	Yes	No				
Are you left or right-handed?	L	R				
If yes to any of the above, please give details:						
Do you consent to being photographed/filmed during this experiment?					Yes	No

I have read the participant guidelines sheet for the above study and had the opportunity to ask questions. I understand that my voluntary participation means I may withdraw at any time without giving any reason. I agree to participate in this study.

Participant Signature

Date.....

7.6.5 Presence Questionnaire

PRESENCE QUESTIONNAIRE

(Witmer & Singer, Vs. 3.0, Nov. 1994)*

Revised by the UQO Cyberpsychology Lab (2004)

Characterize your experience in the environment, by marking an "X" in the appropriate box of the 7-point scale, in accordance with the question content and descriptive labels. Please consider the entire scale when making your responses, as the intermediate levels may apply. Answer the questions independently in the order that they appear. Do not skip questions or return to a previous question to change your answer.

WITH REGARD TO THE EXPERIENCED ENVIRONMENT

1. How much were you able to control events?

|_____|_____|_____|_____|_____|_____|_____|
NOT AT ALL SOMEWHAT COMPLETELY

2. How responsive was the environment to actions that you initiated (or performed)?

|_____|_____|_____|_____|_____|_____|_____|
NOT MODERATELY COMPLETELY
RESPONSIVE RESPONSIVE RESPONSIVE

3. How natural did your interactions with the environment seem?

|_____|_____|_____|_____|_____|_____|_____|
EXTREMELY BORDERLINE COMPLETELY
ARTIFICIAL NATURAL


4. How much did the visual aspects of the environment involve you?

|_____|_____|_____|_____|_____|_____|_____|
NOT AT ALL SOMEWHAT COMPLETELY


5. How natural was the mechanism which controlled movement through the environment?

|_____|_____|_____|_____|_____|_____|_____|
EXTREMELY BORDERLINE COMPLETELY
ARTIFICIAL NATURAL


6. How compelling was your sense of objects moving through space?




7. How much did your experiences in the virtual environment seem consistent with your real world experiences?




8. Were you able to anticipate what would happen next in response to the actions that you performed?




9. How completely were you able to actively survey or search the environment using vision?




10. How compelling was your sense of moving around inside the virtual environment?



11. How closely were you able to examine objects?



12. How well could you examine objects from multiple viewpoints?



13. How involved were you in the virtual environment experience?

|-----|-----|-----|-----|-----|-----|-----|
NOT INVOLVED MILDLY INVOLVED COMPLETELY ENGROSSED

14. How much delay did you experience between your actions and expected outcomes?

|-----|-----|-----|-----|-----|-----|-----|
NO DELAYS MODERATE DELAYS LONG DELAYS

15. How quickly did you adjust to the virtual environment experience?

|-----|-----|-----|-----|-----|-----|-----|
NOT AT ALL SLOWLY LESS THAN

ONE MINUTE

16. How proficient in moving and interacting with the virtual environment did you feel at the end of the experience?

|-----|-----|-----|-----|-----|-----|-----|
NOT PROFICIENT REASONABLY PROFICIENT VERY PROFICIENT

17. How much did the visual display quality interfere or distract you from performing assigned tasks or required activities?

|-----|-----|-----|-----|-----|-----|-----|
NOT AT ALL INTERFERED SOMEWHAT PREVENTED TASK PERFORMANCE

18. How much did the control devices interfere with the performance of assigned tasks or with other activities?

|-----|-----|-----|-----|-----|-----|-----|
NOT AT ALL INTERFERED SOMEWHAT INTERFERED GREATLY

19. How well could you concentrate on the assigned tasks or required activities rather than on the mechanisms used to perform those tasks or activities?

|-----|-----|-----|-----|-----|-----|-----|
NOT AT ALL SOMEWHAT COMPLETELY

IF THE VIRTUAL ENVIRONMENT INCLUDED SOUNDS:

20. How much did the auditory aspects of the environment involve you?

|-----|
NOT AT ALL SOMEWHAT COMPLETELY

21. How well could you identify sounds?

|-----|
NOT AT ALL SOMEWHAT COMPLETELY

22. How well could you localize sounds?

|-----|
NOT AT ALL SOMEWHAT COMPLETELY

IF THE VIRTUAL ENVIRONMENT INCLUDED HAPTIC (SENSE OF TOUCH):

23. How well could you actively survey or search the virtual environment using touch?

|-----|
NOT AT ALL SOMEWHAT COMPLETELY

24. How well could you move or manipulate objects in the virtual environment?

|-----|
NOT AT ALL SOMEWHAT EXTENSIVELY

Last version : March 2013

*Original version : Witmer, B.G. & Singer, M.J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence : Teleoperators and Virtual Environments*, 7(3), 225-240. Revised factor structure: Witmer, B.J., Jerome, C.J., & Singer, M.J. (2005). The factor structure of the Presence Questionnaire. *Presence*, 14(3) 298-312.

7.6.6 Presence Questionnaire Full Results

Question No.	Score 1	Score 2	Score 3	Score 4	Score 5	Median Score	Inter-Quartile Range
1	6	6	6	4	5	6	1
2	6	6	6	4	5	6	1
3	5	5	7	1	5	5	0
4	5	4	6	2	7	5	2
5	7	7	6	4	6	6	1
6	6	6	5	2	4	5	2
7	3	6	5	1	4	4	2
8	4	5	5	2	4	4	1
9	7	6	7	3	5	6	2
10	6	7	7	3	5	6	2
11	7	6	5	3	5	5	1
12	6	6	5	3	5	5	1
13	5	6	4	2	6	5	2
14	4	7	2	4	4	4	0
15	7	7	7	4	5	7	2
16	7	7	6	3	5	6	2
17	3	6	2	3	5	3	2
18	2	2	6	3	5	3	3
19	5	7	7	3	6	6	2
23	5	5	-	2	6	5	1
24	6	5	-	2	4	5	1
Total	112	122	104	58	106	112	7
Realism	37	41	40	15	37	37	3
Possibility to Act	23	23	24	13	19	23	4
Quality of Interface	12	6	11	11	7	11	4
Possibility to Examine	18	19	17	9	16	17	2
Self-evaluation of Performance	14	14	13	7	10	13	4
Haptic	11	10	0	4	10	10	1

7.7 Poster Presented at JVRC 2011

A Virtual Environment for Rugby Skills Training

Helen Miles, Nicholas Musembi, Serban R. Pop, Nigel W. John
School of Computer Science, Bangor University



The Virtual Environment Rugby Skills Trainer (VERST) is an ongoing collaborative project between Bangor University's Schools of Computer Science, Sports, Health and Education Sciences and Psychology, together with Rugby Innovations Ltd., to create an accessible training tool that will help professional rugby players improve certain skills. It is said that it takes 10,000 hours of learning to fully grasp a skill; VERST will reduce the waiting time during practices and allow players to train individually. Although similar projects are being developed for lineouts and defensive strategy, as a first step in our implementation we have chosen to focus on passing the ball; an accurate passing technique is more difficult to achieve than many realize, and is a truly vital base skill. Using a motion capture suite and other tracking hardware, the players' and ball's movements are digitized so to allow them to interact with the virtual environment in a natural way. The accuracy of the environment is preserved by physical modeling and simulating the motion of a virtual ball under real-game conditions. Wales is characterized as a nation that loves rugby, but no one really uses this for marketing – a project from the heartland might be what Wales needs to win the world cup!

Tracking: Real Data Acquisition



Personal Space Tracker
Vicon Motion Capture Suite

Related Work

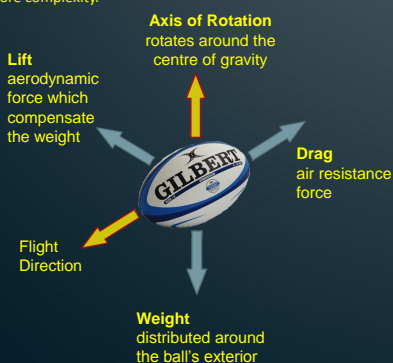
There have been a few projects that use virtual environments for sports applications e.g. [1], but very few that address rugby [2,3]. The improvement in computer graphics, tracking, and display technologies currently provide an opportunity for a high fidelity rugby skills simulator to be created.



Match Rugby Ball with Reflective Markers Applied for Optical Tracking

Physical System of Motion for a Rugby Ball in Flight

- The trajectory is a straight line unless external forces intervene.
- Strictly modelled by aerodynamic forces.
- The initial velocity depends on the launch angle.
- On ascent, weight and drag are acting together, on descent the forces have opposite signs.
- Any external side forces like wind or rain are added to the lift.
- drag or lift can be experimentally determined by throwing the ball at a measured speed.
- rugby ball's surface roughness and the non-symmetrical stitches add more complexity.



Example Simulation Scenario:

- Player covered in reflective markers.
- Wearing stereoscopic glasses looking at projection of virtual world.
- PST captures ball speed and position then sends to program to calculate flight properties.
- Ball tethered to players hands.
- Player throws, virtual ball animated, give statistics and help player understand how to improve technique.

Virtual World

- Multiple targets possible
- Any distance
- Rain, wind can be simulated
- Stereo



Real World



Physics modelled virtual ball follows trajectory the real ball would have

PST captures position and speed of ball as thrown – ball is tethered

Conclusions:

VERST is currently in development but early results augur well for the utility of this tool. As we complete the build of the first prototype, validation tests will be carried out in collaboration with local rugby teams. There are many other skills in this complex sport that we intend to tackle in future work.

References: [1] Vignani, N., Bideau, B., Kulpa, R., Craig, C., Braut, S., & Mullon, F. (2009). Virtual environments for sport analysis: Perception-action coupling in handball goalkeeping. *International Journal of Virtual Reality*, 8(4), 43-48. [2] Braut, S., Bideau, B., Kulpa, R., & Craig, C. (2009). Detecting Deceptive Movement in 1 vs. 1 Based on Global Body Displacement of a Rugby Player. *The International Journal of Virtual Reality*, 8(4), 35-38. [3] Chang, A. K., & Croft, H. (2009). A photogrammetric application in virtual sport training. *The Photogrammetric Record*, 24(125), 51-65.